

Quasi-isometries Between Groups with Two-Ended Splittings

Christopher H. Cashen and Alexandre Martin

January 27, 2016

Abstract

We construct ‘structure invariants’ of a one-ended, finitely presented group that describe the way in which the factors of its JSJ decomposition over two-ended subgroups fit together.

For groups satisfying two technical conditions, these invariants reduce the problem of quasi-isometry classification of such groups to the problem of relative quasi-isometry classification of the factors of their JSJ decompositions. The first condition is that their JSJ decompositions have two-ended cylinder stabilizers. The second is that every factor in their JSJ decompositions is either ‘relatively rigid’ or ‘hanging’. Hyperbolic groups always satisfy the first condition, and it is an open question whether they always satisfy the second.

The same methods also produce invariants that reduce the problem of classification of one-ended hyperbolic groups up to homeomorphism of their Gromov boundaries to the problem of classification of the factors of their JSJ decompositions up to relative boundary homeomorphism type.

1 Introduction

Gromov proposed a program of classifying finitely generated groups up to the geometric equivalence relation of quasi-isometry [20]. A natural approach to this problem is to first try to decompose the group into “smaller” pieces by means of a graph of groups decomposition, and then reduce the quasi-isometry classification problem to the problem of understanding the quasi-isometry types of the various vertex groups and the way these subgroups fit together.

The simplest such decomposition is the decomposition of a group as a graph of groups over finite subgroups. Stallings’s Theorem [42, 43] asserts that a finitely generated group splits as an amalgamated product or an HNN extension over a finite group if and only if it has more than one end. In particular, the existence of a splitting as a graph of groups with finite edge groups is a quasi-isometry invariant. Finitely presented groups admit a maximal splitting over finite subgroups [16], and a theorem of Papasoglu and Whyte [37] says that the collection of quasi-isometry types of one-ended vertex groups of a maximal decomposition of an infinite-ended, finitely presented group is a complete quasi-isometry invariant. This reduces the quasi-isometry classification problem to one-ended groups.

We push this program to its next logical step, which is to decompose one-ended groups over two-ended subgroups. Papasoglu [36] shows that the existence of a splitting of a finitely presented one-ended group over two-ended subgroups is quasi-isometry invariant, provided that the group is not commensurable to a surface group. Moreover, such a group admits a maximal decomposition as a graph of groups over two-ended subgroups, known as a JSJ decomposition, and Papasoglu's results imply that quasi-isometries respect (in a certain sense that will be made precise in Section 2.3.2) the JSJ decomposition. In particular, the quasi-isometry types of the non-elementary vertex groups of the JSJ decomposition are invariant under quasi-isometries. Even more is true: In a non-elementary vertex group of the JSJ we see the collection of conjugates of the two-ended subgroups corresponding to incident edges of the JSJ decomposition. Quasi-isometries of the group must preserve the vertex groups together with their patterns of conjugates of incident edge subgroups. We say a quasi-isometry must preserve the *relative quasi-isometry type* of the vertex. Such patterns have been exploited before [5, 13, 32, 33] to produce quasi-isometry invariants from various 'pattern rigidity' phenomena. However, the quasi-isometry types, or relative quasi-isometry types, of the vertex groups alone do not give complete quasi-isometry invariants, because not all vertex groups have the same coarse intersections: Vertices that are adjacent in the Bass-Serre tree of the decomposition have vertex groups that intersect in two-ended subgroups, while the coarse intersection of non-adjacent vertices may or may not be bounded.

In this paper, we produce further quasi-isometry invariants of a finitely presented one-ended group from an appropriate JSJ decomposition over two-ended subgroups. Under a mild technical restriction, which is known to hold for large classes of groups (see Section 5.2 for a discussion), we show that our invariants give complete quasi-isometry invariants, and thus reduce the quasi-isometry classification problem for such groups to relative versions of these problems in the vertex groups of their JSJ decomposition.

In the case of one-ended hyperbolic groups, a weaker classification is possible, namely the classification up to homeomorphisms of Gromov boundaries. Indeed, recall that quasi-isometric hyperbolic groups have homeomorphic Gromov boundaries at infinity, but there are examples of hyperbolic groups with homeomorphic boundary that are not quasi-isometric. For a hyperbolic group, the existence of a splitting over a finite subgroup amounts to having a disconnected Gromov boundary, and is thus detected by the homeomorphism type of the boundary. Paralleling the results of Papasoglu and Whyte for quasi-isometries, Martin and Świątkowski [31] show that hyperbolic groups with infinitely many ends have homeomorphic boundaries if and only if they have the same sets of homeomorphism types of boundaries of one-ended factors, reducing the classification problem to the case of one-ended hyperbolic groups.

For a one-ended hyperbolic group G whose boundary is not a circle, Bowditch shows that there exists a splitting over a two-ended subgroup if and only if there exists a cut pair in the boundary [9]. From the structure of cut pairs in the boundary, he deduces the existence of a simplicial tree on which $\text{Homeo}(\partial G)$ acts by isomorphisms. Paralleling the quasi-isometry case, the action of $\text{Homeo}(\partial G)$ on this tree preserves the

relative boundary homeomorphism types of vertex stabilizers in G .

Using the same approach as for our quasi-isometry classification, we construct a complete system of invariants of the homeomorphism type of the Gromov boundary of a hyperbolic group, which completely reduces the classification problem for hyperbolic groups to relative versions of this problem in the vertex groups of their JSJ decomposition.

Our results rely heavily on the existence of a canonical choice of tree that is preserved by quasi-isometries, respectively, homeomorphisms of the boundary, with the additional property that any such map induces maps of the same nature at the level of the vertex groups.

In the case of hyperbolic groups, the tree is Bowditch's canonical JSJ tree constructed from cut pairs in the boundary. For a more general finitely presented one-ended group G , let $\mathbf{\Gamma}$ be a JSJ decomposition over two-ended subgroups. Let $T := T(\mathbf{\Gamma})$ be the Bass-Serre tree of $\mathbf{\Gamma}$. Commensurability of stabilizers defines an equivalence relation on edges of T whose equivalence classes are called *cylinders*. Guirardel and Levitt [23] show that the dual tree to the covering of T by cylinders defines a new tree $\text{Cyl}(T)$, the *tree of cylinders* of T , with a cocompact G -action, and, in fact, this tree is independent of $\mathbf{\Gamma}$, so it makes sense to call it the *tree of cylinders of G* , and denote it $\text{Cyl}(G)$. It follows from Papasoglu's results, see Theorem 2.9, that a quasi-isometry between finitely presented one-ended groups induces an isomorphism between their trees of cylinders, and restricts to give a quasi-isometry of each vertex group G_v of the tree of cylinders that coarsely preserves the pattern \mathcal{P}_v of edge groups G_e , for edges e incident to v . When the edge stabilizer of $\text{Cyl}(G)$ in G are two-ended then the quotient graph of groups gives a canonical decomposition of G over two-ended subgroups. Our results are strongest when additionally the cylinder stabilizers are two-ended. This is the case, in particular, when G is hyperbolic, in which case $\text{Cyl}(G)$ is equivariantly isomorphic to Bowditch's tree.

Before giving more details about the ideas behind our classifications, we restrict to the case that G is a one-ended hyperbolic group, and we adopt some notation that will allow us to discuss simultaneously the quasi-isometry and boundary homeomorphism cases. Let $\text{Map}((X, \mathcal{P}^X), (Y, \mathcal{P}^Y))$ denote alternately:

- The set of quasi-isometries from X to Y taking the collection of coarse equivalence classes of subsets \mathcal{P}^X to the collection of coarse equivalence classes of subsets \mathcal{P}^Y .
- The set of boundary homeomorphisms from ∂X to ∂Y taking the collection of subsets $\partial\mathcal{P}^X$ to the collection of subsets $\partial\mathcal{P}^Y$.

Similarly, $\text{Map}(G)$ is either $\mathcal{QI}(G)$ or $\text{Homeo}(\partial G)$. An element of $\text{Map}(\dots)$ will be referred to as a Map-equivalence.

The idea behind our classification result is the following. First remark that a finitely presented one-ended group is quasi-isometric to a complex of spaces over the JSJ tree of cylinders (this will be recalled in Section 2.4). Moreover, such a decomposition as a complex of spaces is compatible with Map-equivalences, in that a Map-equivalence

between two finitely presented one-ended groups coarsely preserves the structure of complex of spaces, as follows from the aforementioned results of Papasoglu and Bowditch. To determine whether two groups are Map-equivalent, we thus want to decide whether their JSJ trees of cylinders are isomorphic and then try to promote such an isomorphism to a Map-equivalence between the complexes of spaces, hence, between the groups. In full generality, a Map-equivalence not only induces a simplicial isomorphism between the trees of cylinders, but also preserves extra information about the vertex groups that we want to take into account: Map-class of vertex groups, relative Map-class with respect to incident edge groups, etc. In other words, we decorate the vertices of the trees of cylinders with these additional pieces of information. What we then want is to understand when there exists a decoration-preserving isomorphism between the tree of cylinders of two groups. We want to add enough additional information so that a decoration-preserving isomorphism between the trees of cylinders can be promoted to a Map-equivalence between the associated groups. Moreover, we would like to have an algorithm telling us when two such decorated trees of cylinders are isomorphic.

This latter point will be dealt with by generalizing to cocompact decorated trees a theorem from graph theory giving a necessary and sufficient condition for the universal covers of two graphs to be isomorphic [27], see Section 4.2. To such a decorated tree we will associate a *structure invariant* that completely determines the tree up to decoration-preserving isomorphism.

To now get an intuition of the decorations we consider in this article, let us start from the decoration of the tree of cylinders that associates to each vertex v the relative Map-equivalence $\text{Map}(G_v, \mathcal{P}_v)$ mentioned earlier, where \mathcal{P}_v is the peripheral structure coming from the incident edge groups, and let us try to promote a decoration-preserving isomorphism χ between trees of cylinders to a Map-equivalence between groups. The goal is to choose, for each vertex v , a Map-equivalence between v and $\chi(v)$, and piece them together to get Map-equivalence between groups.

The first potential problem is a realization problem. We know that v and $\chi(v)$ are Map-equivalent, but we also need to know that there is such a Map-equivalence that matches up peripheral subsets in the same way that χ matches up edges incident to v and $\chi(v)$.

The second potential problem is that the vertex Map-equivalences must agree when their domains overlap. For boundary homeomorphisms the overlap is just the boundary of an edge space, which is a pair of points, so this a matter of choosing consistent *orientations* on the edge spaces.

Roughly, this is the phenomenon underlying the fact that two HNN extensions of the form

$$G_+ := \langle G, t | t u \bar{t} = v \rangle \quad \text{and} \quad G_- := \langle G, t | t u \bar{t} = \bar{v} \rangle$$

might not be Map-equivalent in general, for a group G and elements u, v of G .

It is worth noting that the orientation obstruction automatically disappears if all the edge groups contain an infinite dihedral group.

For the classification up to boundary homeomorphisms, these three obstructions,

relative type, realizability, and orientation, are essentially the only obstructions to constructing a Map-equivalence between the groups. We use these considerations to produce a finer decoration that yields a structure invariant that is a complete invariant for boundary homeomorphism type, see Theorem 7.1.

In the case of quasi-isometries, orientation of the edge spaces is not enough; we must choose Map-equivalence of the vertex spaces that agree all along the length of shared edge spaces. There are two cases in which we can decide if this is possible. The first is that vertex spaces are extremely flexible so that we have a lot of freedom to choose Map-equivalences and make them agree on edge spaces. This is the case for the so-called ‘hanging’ vertices. The other case is the opposite one, in which the vertex space is extremely ‘rigid’ and we have very little choice about how to choose the maps. In this case we define an invariant called a *stretch factor* that we then incorporate into the decoration. If all vertices are either rigid or hanging then the decoration that takes into account vertex relative quasi-isometry type, realizability, orientation, and stretch factors yields a complete quasi-isometry invariant, see Theorem 8.5. This notion of *relative quasi-isometric rigidity* is known to hold for many classes of groups, see Section 5.2.

The classifications we provide are inherently more technical than the classifications for splittings over finite subgroups by Papasoglu and Whyte and Martin and Świątkowski. Underlying those theorems is the fact that in splittings over finite groups the vertex groups of the canonical decomposition are essentially independent of one another, so only the pieces of the decomposition matter. Our classification must handle not just the pieces, but also their complex interactions.

1.1 Applications

Our main results give invariants that reduce classification problems for finitely presented, one-ended groups to a relative version of the classification problem on the vertex groups of a JSJ decomposition.

Note that vertex groups of a two-ended JSJ decomposition of a one-ended group are not necessarily one-ended, but we really have achieved a significant reduction of the classification problem because the problem of classification of relative quasi-isometry types is far more constrained than the general problem. The most striking example of this phenomenon is a theorem of Cashen and Macura [14] that says that if a free group appears as a vertex group of a JSJ decomposition then it is either hanging or it is quasi-isometrically rigid relative to the peripheral structure induced by incident edge groups. Moreover, in the rigid case the relative quasi-isometry question reduces to a relative *isometry* question in certain associated cube complexes, which seems far more tractable.

A sample application of our main results is a complete description, in terms of [14], of the quasi-isometry and boundary homeomorphism types of one-ended hyperbolic groups that split as graphs of groups¹ with free vertex groups and cyclic edge groups.

Consider the following example:

¹ Such a decomposition can be algorithmically improved to a JSJ decomposition of the same type [12].

Example 1.1. Let $G_i = \langle a, b, t \mid \bar{t}u_i t = v_i \rangle$, where u_i and v_i are words in $\langle a, b \rangle$ given below. In each case G_i should be thought of as an HNN extension of $\langle a, b \rangle$ over \mathbb{Z} with stable letter t .

Let $u_0 := a$, $v_0 := ab\bar{a}\bar{b}^2$, $u_1 := ab$, $v_1 := a^2\bar{b}^2$, $u_2 := ab^2$, and $v_2 := a^2\bar{b}$.

Then G_0 , G_1 , and G_2 are pairwise non-quasi-isometric, but have homeomorphic Gromov boundaries.

The quasi-isometry claim follows because the vertex groups in each case are quasi-isometrically rigid relative to the incident edge groups, and the stretch factor for G_i is the ratio of the word lengths of u_i and v_i , which are different for each i ; see Section 3 and Section 5.7. The groups have homeomorphic boundaries because the stretch factors are the only difference, and boundary homeomorphism is not sensitive to stretch factors; see Section 3 and Example 7.2.

Our methods produce interesting invariants even in cases that the groups are not hyperbolic or that the relative problems for the vertex groups is not be completely understood:

Example 1.2. Let M be the mapping class group of a non-sporadic hyperbolic surface, with a fixed finite generating set. Let $g_0, g_1, g'_0, g'_1 \in M$ be pseudo-Anosov elements that are not proper powers. Let $G := M *_\mathbb{Z} M$ and $G' := M *_\mathbb{Z} M$ be the amalgamated product groups obtained by identifying g_0 with g_1 and g'_0 with g'_1 , respectively. Then G and G' are not quasi-isometric if the ratio of translation lengths of g_0 and g_1 is different, up to inversion, from the ratio of translation lengths of g'_0 and g'_1 .

This follows because mapping class groups are quasi-isometrically rigid and the ratios of the translation lengths in this example are the stretch factors; see Section 5.2 and Section 5.

Acknowledgements

The first author thanks Mladen Bestvina, Jason Behrstock, and Gilbert Levitt for interesting conversations related to this work. In particular, the idea that a version of Leighton's Theorem would provide a concise description of our structure invariants was suggested by Bestvina, and Example 1.2 was suggested by Behrstock.

Both authors were supported by the European Research Council (ERC) grant of Goulmira Arzhantseva, grant agreement #259527 and the Erwin Schrödinger Institute workshop "Geometry of Computation in Groups". The first author is partially supported by the Austrian Science Fund (FWF):M1717-N25. The second author is partially supported by the Austrian Science Fund (FWF):M1810-N25.

2 Preliminaries

We assume familiarity with standard concepts such as Cayley graphs, ends of spaces, and (Gromov) hyperbolic geometry. See [10] for background.

A group is *virtually cyclic* if it has an infinite cyclic subgroup of finite index. A group is *non-elementary* if it is neither finite nor virtually cyclic. A standard exercise is to show that a finitely generated, two-ended group is virtually cyclic.

2.1 Coarse geometry

Throughout the paper the qualifier ‘coarse’ is used to indicate ‘up to additive error’.

Let (X, d_X) and (Y, d_Y) be metric spaces. Subsets of X are *coarsely equivalent* if they are bounded Hausdorff distance from one another. A subset A is *coarsely contained* in B if A is coarsely equivalent to a subset of B . Two maps ϕ and ϕ' from X to Y are said to be *coarsely equivalent* or *bounded distance* from each other if $\sup_{x \in X} d_Y(\phi(x), \phi'(x)) < \infty$. A map $\phi: X \rightarrow Y$ is *coarsely surjective* if its image is coarsely equivalent to Y .

A map $\phi: X \rightarrow Y$ is a *controlled embedding*² if there exist unbounded, non-decreasing real functions ρ_0 and ρ_1 such that for all x and x' in X we have:

$$\rho_0(d_X(x, x')) \leq d_Y(\phi(x), \phi(x')) \leq \rho_1(d_X(x, x'))$$

There are several classes of controlled embeddings that have special names. If $\rho_i(r) := M_i r$ for $i \in \{0, 1\}$ then ϕ is

- an *isometric embedding* if $M_0 = 1 = M_1$,
- an *M-similitude* if $M_0 = M = M_1$,
- an *M-biLipschitz embedding* if $M_0^{-1} = M = M_1$.

If such a map is surjective then it is called, respectively, an *isometry*, *similarity*, or *biLipschitz equivalence*.

For each of these, we can add the qualifier ‘coarse’ and allow an additive error, so that $\rho_0(r) := M_0 r - A$ and $\rho_1(r) := M_1 r + A$. For instance, ϕ is an (M, A) -coarse biLipschitz embedding if ρ_i is as above with $M_0 = M^{-1}$ and $M_1 = M$. In the ‘coarse’ cases we drop the term ‘embedding’ if the map is coarsely surjective.

An (M, A) -coarse biLipschitz embedding is more commonly called an (M, A) -*quasi-isometric embedding*.

Let $\text{QIsom}(X, Y)$ denote the set of quasi-isometries from X to Y , and let $\text{CIsom}(X, Y)$ denote the set of coarse isometries from X to Y . For a coarsely surjective map $\phi: X \rightarrow Y$, a *coarse inverse* is a map $\bar{\phi}: Y \rightarrow X$ such that $\bar{\phi} \circ \phi$ is coarsely equivalent to Id_X and $\phi \circ \bar{\phi}$ is coarsely equivalent to Id_Y . If $\phi \in \text{QIsom}(X, Y)$ then all coarse inverses of ϕ are coarsely equivalent and belong to $\text{QIsom}(Y, X)$. If $\phi \in \text{CIsom}(X, Y)$ then every coarse inverse of ϕ belongs to $\text{CIsom}(Y, X)$.

Let $\mathcal{I}(X, Y)$, $\mathcal{CI}(X, Y)$, and $\mathcal{QI}(X, Y)$ denote, respectively, the sets $\text{Isom}(X, Y)$, $\text{CIsom}(X, Y)$, and $\text{QIsom}(X, Y)$, modulo coarse equivalence. When $Y = X$ we shorten the notation to $\mathcal{I}(X)$, $\mathcal{CI}(X)$, $\mathcal{QI}(X)$, and each of these form a group under composition.

²This is more commonly called a ‘coarse embedding’.

A subset of $\mathcal{CI}(X, Y)$ or $\mathcal{QI}(X, Y)$ is said to be *uniform* if there exists a C such that every element of the subset is an equivalence class of maps containing a C -coarse isometry or a (C, C) -quasi-isometry, respectively.

Quasi-isometries respect coarse equivalence of subsets. If \mathcal{P} is a set of coarse equivalence classes of subsets of X , and \mathcal{P}' is a set of coarse equivalence classes of subsets of Y , let $\mathcal{QI}((X, \mathcal{P}), (Y, \mathcal{P}'))$ be the subset of $\mathcal{QI}(X, Y)$ consisting of quasi-isometries that induce bijections between \mathcal{P} and \mathcal{P}' . Similarly, $\mathcal{QI}((X, \mathcal{P})) := \mathcal{QI}((X, \mathcal{P}), (X, \mathcal{P}))$ is a subgroup of $\mathcal{QI}(X)$.

If $\phi: X \rightarrow Y$ is a quasi-isometry, define $\phi_*: \mathcal{QI}(X) \rightarrow \mathcal{QI}(Y)$ by $\phi_*(\psi) := \phi \circ \psi \circ \bar{\phi}$.

A subset $Y \subset X$ is *coarsely connected* if there exists some A such that for every y and $y' \in Y$ there is, for some k , a chain $y_0 := y, y_1, \dots, y_k := y'$ of $(k+1)$ points $y_i \in Y$ such that $d(y_i, y_{i+1}) \leq A$.

A *geodesic* is an isometric embedding of a connected subset of \mathbb{R} . A *coarse geodesic* is a coarse geodesic embedding of a coarsely connected subset of \mathbb{R} . A *quasi-geodesic* is a quasi-isometric embedding of a coarsely connected subset of \mathbb{R} .

The space X is said to be *geodesic*, *A-coarse geodesic*, or *(M, A)-quasi-geodesic* if for every pair of points in X there exists, respectively, a geodesic, *A-coarse geodesic*, or *(M, A)-quasi-geodesic* connecting them.

Let $\llbracket X \rrbracket$ denote the set of proper geodesic metric spaces quasi-isometric to X . If \mathcal{P} is a set of coarse equivalence classes of subsets of X , let $\llbracket (X, \mathcal{P}) \rrbracket$ denote the set of pairs (Y, \mathcal{P}') where Y is a geodesic metric space and \mathcal{P}' is a collection of coarse equivalence classes of subsets of Y such that there exists a quasi-isometry from X to Y that induces a bijection from \mathcal{P} to \mathcal{P}' . We call $\llbracket X \rrbracket$ the *quasi-isometry type* of X and $\llbracket (X, \mathcal{P}) \rrbracket$ the *relative quasi-isometry type* of (X, \mathcal{P}) .

If L is a path connected subset of a geodesic metric space (X, d_X) , let d_L denote the induced length metric on L . A *quasi-line* in X is a path connected subset L such that (L, d_L) is quasi-isometric to \mathbb{R} and the inclusion map of L into X is a controlled embedding.

We define a *peripheral structure* \mathcal{P} on geodesic metric space X to be a collection of coarse equivalence classes of quasi-lines. In particular, if G is a finitely generated group and \mathcal{H} is a finite collection of two-ended subgroups of G , then \mathcal{H} induces a peripheral structure consisting of distinct coarse equivalence classes of conjugates of elements of \mathcal{H} .

2.2 Graphs of groups

Let Γ be a finite oriented graph. Let $\mathcal{V}\Gamma$ be the set of vertices of Γ , and let $\mathcal{E}^+\Gamma$ be the set of oriented edges. For $e \in \mathcal{E}^+\Gamma$, let $\iota(e)$ be its initial vertex, and let $\tau(e)$ be its terminal vertex. For each $e \in \mathcal{E}^+\Gamma$ formally define \bar{e} to be an inverse edge to e with $\iota(\bar{e}) := \tau(e)$, $\tau(\bar{e}) := \iota(e)$, and $\bar{\bar{e}} := e$. The inverse edge \bar{e} should be thought of as e traversed against its given orientation. Let $\mathcal{E}^-\Gamma$ denote the set of inverse edges, and $\mathcal{E}\Gamma := \mathcal{E}^+\Gamma \cup \mathcal{E}^-\Gamma$.

A graph of groups $\mathbf{\Gamma} := (\Gamma, \{G_\gamma\}_{\gamma \in \mathcal{V}\Gamma \cup \mathcal{E}^+\Gamma}, \{\varepsilon_e\}_{e \in \mathcal{E}^+\Gamma})$ consists of a finite directed graph Γ , groups G_γ for $\gamma \in \mathcal{V}\Gamma \cup \mathcal{E}^+\Gamma$ such that $G_e < G_{\iota(e)}$ for $e \in \mathcal{E}^+\Gamma$, and injections $\varepsilon_e: G_e \hookrightarrow G_{\tau(e)}$ for $e \in \mathcal{E}^+\Gamma$.

For symmetry in the notation it is convenient to define $G_{\bar{e}} := G_e$ for each $\bar{e} \in \mathcal{E}^- \Gamma$ and let $\varepsilon_{\bar{e}}$ denote the inclusion of $G_{\bar{e}} := G_e$ into $G_{\tau(\bar{e})} := G_{\iota(e)}$.

A graph of groups Γ has an associated fundamental group $G = G(\Gamma)$ obtained by amalgamating the vertex groups over the edge groups [41]. We say that Γ is a *graph of groups decomposition* of G .

The Bass-Serre tree $T := T(\Gamma)$ of Γ is the tree on which G acts without edge inversions, such that $G \backslash T = \Gamma$ and such that the stabilizer G_t of $t \in \mathcal{V}T \cup \mathcal{E}T$ is a conjugate in G of the group $G_{\mathbf{t}}$, where \mathbf{t} is the image of t under the quotient map $T \rightarrow \Gamma$.

Throughout we use the notation \mathbf{t} to denote the image of t in Γ . Conversely, for each $\gamma \in \mathcal{V}\Gamma \cup \mathcal{E}\Gamma$ we choose some lift $\tilde{\gamma}$ of γ to T . Given a maximal subtree in Γ we can, and do, choose lifts of vertices and edge in the subtree to get a subtree in T .

Definition 2.1. Given a vertex group G_v of Γ , the *peripheral structure coming from incident edge groups*, \mathcal{P}_v , is the set of distinct coarse equivalence classes in G_v of G_v -conjugates of the images of the edge injections $\varepsilon_e: G_e \hookrightarrow G_v$ for edges $e \in \mathcal{E}\Gamma$ with $\tau(e) = v$.

Definition 2.2. Given a vertex v of $T(\Gamma)$, the *peripheral structure coming from incident edge groups*, \mathcal{P}_v , is the set of distinct coarse equivalence classes in the stabilizer subgroup G_v of v of stabilizers of incident edge groups.

It is immediate from the definitions that the quotient by the G -action identifies (G_v, \mathcal{P}_v) and $(G_{\mathbf{v}}, \mathcal{P}_{\mathbf{v}})$.

We are interested in graphs of groups in which the edge groups are two-ended, hence virtually cyclic. Commensurability of edge stabilizers defines an equivalence relation on the edges of the Bass-Serre tree T of such a splitting. The equivalence classes of edges are called *cylinders*. Every cylinder is a subtree of T [23, Lemma 4.2]. It follows that we get another tree $\text{Cyl}(T)$, called the *tree of cylinders* of T , by taking the dual tree to the covering of T by cylinders.

Definition 2.3. A graph of groups with two-ended edge groups is *k-acylindrical* if the cylinders of its Bass-Serre tree have diameter at most k . It is *acylindrical* if there exists k such that it is k -acylindrical.

Let C be a cylinder in T and let $\text{Stab}(C)$ be the stabilizer of C in G . Choose an infinite order element z in G_e for some edge $e \in C$. For any element $g \in \text{Stab}(C)$, G_e and G_{ge} are commensurable, virtually cyclic groups. Since $\langle z \rangle$ is a finite index subgroup of G_e , there exist non-zero a and b such that $gz^a g^{-1} = z^b$. Define $\Delta(g) = \frac{a}{b}$. This defines a homomorphism $\Delta: \text{Stab}(C) \rightarrow \mathbb{Q}^*$, called the *modular homomorphism* of C , that is independent of the choice of z .

Definition 2.4. A cylinder is called *unimodular* if the image of its modular homomorphism is in $\{-1, 1\}$.

A graph of groups with two-ended edge groups is *unimodular* if all of its cylinders are unimodular.

We also define a *modulus* $\hat{\Delta}$ on pairs of edges of a cylinder C :

Definition 2.5. Let e_0 and e_1 be edges in C . Let $\langle z_0 \rangle < G_{e_0}$ and $\langle z_1 \rangle < G_{e_1}$ be infinite cyclic subgroups of minimal index. Define $\hat{\Delta}(e_0, e_1) = \frac{[\langle z_1 \rangle : \langle z_0 \rangle \cap \langle z_1 \rangle]}{[\langle z_0 \rangle : \langle z_0 \rangle \cap \langle z_1 \rangle]}$.

It is easy to check that $\hat{\Delta}$ does not depend on the choice of minimal index infinite cyclic subgroups, and that $\hat{\Delta}(e, ge) = |\Delta(g)|$.

2.3 The JSJ tree of cylinders

2.3.1 Definitions

A JSJ decomposition of a finitely presented, one-ended group G over two-ended subgroups is a graph of groups with two-ended edge groups that encodes all splittings of G over two-ended subgroups. Equivalent descriptions of such decompositions appear in Dunwoody and Sageev [17], Fujiwara and Papasoglu [19], and Guirardel and Levitt [22]. See also Rips and Sela [39].

Following Papasoglu [36], we will give a geometric description of JSJ decompositions. First, we need some terminology.

A quasi-line L is *separating* if its complement has at least two *essential* components, that is, components that are not contained in any finite neighborhood of L . In particular, if G splits over a two-ended subgroup then that two-ended subgroup is bounded distance from a separating quasi-line. Separating quasi-lines *cross* if each travels arbitrarily deeply into two different essential complementary components of the other.

Let G_v be a vertex group in a graph of groups decomposition. Let \mathcal{P}_v be the peripheral structure on G_v coming from incident edge groups. Let Σ be a hyperbolic pair of pants. Let $\mathcal{P}_{\partial\Sigma}$ be the peripheral structure on the universal cover $\tilde{\Sigma}$ of Σ consisting of the coarse equivalence classes of the components of the preimages of the boundary curves.

Definition 2.6. We say v is *hanging*³ if (G_v, \mathcal{P}_v) is quasi-isometric to $(\tilde{\Sigma}, \mathcal{P}_{\partial\Sigma})$. We say v is *rigid* if it is not two-ended, not hanging, and does not split over a two-ended subgroup relative to its incident edge groups.

The rigid and hanging terminology extends to vertices in $T(\mathbf{\Gamma})$ in the obvious way.

Definition 2.7. Let G be a finitely presented one-ended group that is not commensurable to a surface group. A *JSJ decomposition* of G is a (possibly trivial) graph of groups decomposition $\mathbf{\Gamma}$ with two-ended edge groups satisfying the following conditions:

- (a) Every vertex group is either two-ended, hanging, or rigid.
- (b) If v is a valence one vertex with two-ended vertex group then the incident edge group does not surject onto G_v .
- (c) Every cylinder in the Bass-Serre tree of $\mathbf{\Gamma}$ that contains exactly two hanging vertices also contains a rigid vertex.

³After the ‘quadratically hanging’ vertex groups of Rips and Sela [38].

This definition is equivalent to those cited above. The essential facts are that:

1. Hanging vertices contain crossing pairs of separating quasi-lines.
2. Every pair of crossing separating quasi-lines is coarsely contained in a conjugate of a hanging vertex group.
3. A separating quasi-line that is not crossed by any other separating quasi-line is coarsely equivalent to a conjugate of an edge group.
4. Every edge group is coarsely equivalent to a separating quasi-line that is not crossed by any other separating quasi-line.

Remark. Condition (c) implies that the hanging vertex groups are maximal hanging, which is necessary for item 4.

Remark. The case that a vertex group is the fundamental group of a pair of pants and the incident edge groups glue on to the boundary curves is called ‘rigid’ in the usual JSJ terminology because there are no splittings of the pair of pants group relative to the boundary subgroups. Algebraically, such a vertex behaves like our rigid vertices, but geometrically this is a hanging vertex.

In general a group does not have a unique JSJ decomposition, but rather a deformation space of JSJ decompositions [18, 22]. Furthermore, all JSJ decompositions are in the same deformation space, which means any one can be transformed into any other by means of a finite sequence of moves of a prescribed type. The tree of cylinders of a decomposition depends only on the deformation space [23, Theorem 1], up to G -equivariant isomorphism, so there is a unique JSJ tree of cylinders.

Definition 2.8. Let G be a finitely presented one-ended group. The *JSJ tree of cylinders* $\text{Cyl}(G)$ of G is the tree of cylinders of the Bass-Serre tree of any JSJ decomposition of G over two-ended subgroups.

The quotient graph of groups $G \backslash \text{Cyl}(G)$ gives a canonical decomposition of G . It is canonical in the sense that its Bass-Serre tree is G -equivariantly isomorphic to the tree of cylinders of any JSJ decomposition of G over two-ended subgroups. However, such a graph of cylinders is *not necessarily* a JSJ decomposition, and it does not even have two-ended edge groups, in general. We return to this issue in Section 2.3.4.

2.3.2 Quasi-isometry invariance of the JSJ tree of cylinders

Since quasi-isometries coarsely preserve quasi-lines, and preserve the crossing and separating properties of quasi-lines, the following version of quasi-isometry invariance of JSJ decompositions follows from Papasoglu’s work:

Theorem 2.9 (cf [36, Theorem 7.1]). *Let G and G' be finitely presented one-ended groups. Suppose $\phi: G \rightarrow G'$ is a quasi-isometry. Then there is a constant C such that ϕ induces an isomorphism $\phi_*: \text{Cyl}(G) \rightarrow \text{Cyl}(G')$ that preserves vertex type — cylinder, hanging, or rigid — and for $v \in \mathcal{V} \text{Cyl}(G)$ takes G_v to within distance C of $G'_{\phi_*(v)}$.*

Proof. Two-ended subgroups of G are bounded distance from each other if and only if they are commensurable, so cylinders are coarse equivalence classes of edge and two-ended vertex stabilizers. These are exactly the coarse equivalence classes of separating quasi-lines that are not crossed by other separating quasi-lines, so quasi-isometries induce a bijection between cylinders. The remaining statements are from Papasoglu [36, Theorem 7.1] for the rigid and hanging vertex stabilizers and Vavrichek [45] for the cylinder stabilizers. \square

Corollary 2.10. *If v is a rigid or hanging vertex in $\text{Cyl}(G)$ then*

$$\phi_v := \pi_{\phi_*(v)} \circ \phi|_{G_v} \in \mathcal{QI}((G_v, \mathcal{P}_v), (G'_{\phi_*(v)}, \mathcal{P}_{\phi_*(v)}))$$

where $\pi_{\phi_*(v)}$ takes the image of $\phi|_{G_v}$ to $G'_{\phi_*(v)}$ by closest point projection.

Remark. $\pi_{\phi_*(v)}$ is coarsely well defined since $\phi(G_v)$ is within distance C of $G'_{\phi_*(v)}$.

2.3.3 Boundary homeomorphism invariance of the JSJ tree of cylinders

In the case of a one-ended hyperbolic group that is not cocompact Fuchsian, Bowditch constructed a canonical JSJ splitting of the group directly from the combinatorics of the local cut points of the Gromov boundary of the group [9]. In this case, he proves that the JSJ tree is unique, and such a tree is thus equivariantly isomorphic to the JSJ tree of cylinders of the group. As a homeomorphism between the Gromov boundaries of two hyperbolic groups preserves the topology, we get the following:

Theorem 2.11 (cf [9]). *Let G and G' be one-ended hyperbolic groups that are not cocompact Fuchsian. Suppose $\rho: \partial G \rightarrow \partial G'$ is a homeomorphism between their Gromov boundaries. Then ρ induces an isomorphism $\rho_*: \text{Cyl}(G) \rightarrow \text{Cyl}(G')$ that preserves vertex type — cylinder, hanging, or rigid — and for $v \in \mathcal{V} \text{Cyl}(G)$ the homeomorphism ρ restricts to a homeomorphism $\rho|_{\partial G_v}: \partial G_v \rightarrow \partial G'_{\rho_*(v)}$.*

Corollary 2.12. *For every vertex $v \in \text{Cyl}(G)$:*

$$\rho_v := \rho|_{\partial G_v} \in \text{Homeo}((\partial G_v, \partial \mathcal{P}_v), (\partial G'_{\rho_*(v)}, \partial \mathcal{P}'_{\rho_*(v)}))$$

2.3.4 Improved invariants from restrictions on the JSJ tree of cylinders

The JSJ tree of cylinders $\text{Cyl}(G)$ of a finitely presented one-ended group G suffices for the definition of the basic quasi-isometry invariants of Section 4. These are far from complete invariants, however.

We can refine the invariants in restricted classes of groups. For instance, if the edges of $\text{Cyl}(G)$ have two-ended stabilizers in G then we can define stretch factors as in Section 5.

When the cylinder stabilizers are two-ended then the full power of Section 6 can be brought to bear. In this case the graph of cylinders is a canonical JSJ decomposition of G over two-ended subgroups. This is the case of chief interest for this paper, and this is always the case if G is hyperbolic.

If the cylindrical vertices of $\text{Cyl}(G)$ have two-ended stabilizers then they are all finite valence in $\text{Cyl}(G)$. Furthermore, if $\mathbf{\Gamma}$ is a JSJ decomposition of G over two-ended subgroups and $v \in T(\mathbf{\Gamma})$ is a vertex whose stabilizer is rigid or hanging then v belongs to more than one cylinder, so $\text{Cyl}(G)$ has a vertex corresponding to v with the same stabilizer subgroup in G . Thus, $\text{Cyl}(G)$ is bipartite, with one part, V_C , consisting of finite valence cylindrical vertices, one for each cylinder of $T(\mathbf{\Gamma})$, and the other part consisting of the sets of vertices V_H and V_R , which are all of infinite valence, corresponding to hanging and rigid vertices of $T(\mathbf{\Gamma})$, respectively.

See [23, Proposition 5.2] for a general result about when the tree of cylinders gives a JSJ decomposition.

2.4 Trees of spaces

Let T be an oriented simplicial tree. For each vertex $v \in T$ let X_v be a metric space. For each edge $e \in \mathcal{E}^+T$ let X_e be a subspace of $X_{\iota(e)}$, and let $\alpha_e: X_e \rightarrow X_{\tau(e)}$ be a map such that for $X_{\bar{e}} := \alpha_e(X_e)$ there exists a map $\alpha_{\bar{e}}: X_{\bar{e}} \rightarrow X_{\iota(e)}$ such that $\alpha_{\bar{e}} \circ \alpha_e$ is bounded distance from Id_{X_e} and $\alpha_e \circ \alpha_{\bar{e}}$ is bounded distance from $\text{Id}_{X_{\bar{e}}}$.

Let X be the quotient of the set

$$\coprod_{v \in \mathcal{V}T} X_v \sqcup \coprod_{e \in \mathcal{E}^+T} \coprod_{x \in X_e} \{x\} \times e$$

by the identifications $x \sim (x, \iota(e))$ and $\alpha_e(x) \sim (x, \tau(e))$. We call

$$X := X(T, \{X_t\}_{t \in \mathcal{V}T \cup \mathcal{E}T}, \{\alpha_e\}_{e \in \mathcal{E}T})$$

a *tree of spaces over T* . The X_v are called *vertex spaces* and the X_e are called *edge spaces*. The sets $\{x\} \times e$ we call *rungs*, and metrize them as unit intervals. The maps α_e are called *attaching maps*.

We say X has *locally finite edge patterns* if for every vertex v , every $x \in X_v$, and every $R \geq 0$, there are finitely many $e \in \mathcal{E}T$ such that X_e intersects $\bar{\mathcal{B}}_R^v(x) := \{y \in X_v \mid d_{X_v}(x, y) \leq R\}$.

A *k-chain* is a sequence of points $x_0, x'_0, x_1, \dots, x'_k$ such that x_i and x'_i are contained in a common vertex space or rung for all $0 \leq i \leq k$ and x'_i and x_{i+1} are contained in a common vertex space or rung for $0 \leq i < k$. The *length* of a k -chain is $d(x_0, x'_0) + \sum_{0 < i \leq k} (d(x'_{i-1}, x_i) + d(x_i, x'_i))$, where the distance terms are interpreted in the vertex space or rung containing the pair. The quotient pseudo-metric on X defines the distance between two points to be the infimum of lengths of chains joining them. When x and x' are points in vertex spaces then for a k -chain joining them we have that x'_i and x_{i+1} are the endpoints of a rung for each $0 \leq i < k$, so the length of the chain is at least k . The following properties follow easily from this observation.

Lemma 2.13. *For a tree of spaces X :*

- *The quotient pseudo-metric is a metric.*

- A ball of radius at most 1 in a vertex space is isometrically embedded in X .
- The rungs are isometrically embedded.

Lemma 2.14. *If X is a tree of spaces over T such that each vertex space is proper and geodesic, each edge space is closed and discrete, and edge patterns are locally finite in each vertex space, then X is a proper geodesic space.*

Proof. For properness it suffices to show a closed ball centered at a point in a vertex space is compact, since a ball centered at an interior point in a rung is contained in the union of balls of radius 1 larger centered at the endpoints of the rung.

Take $R > 0$ and a point $x \in X_v$ for some vertex v . Let $N_0 := \overline{B}_R^v(x)$. Since edge patterns are locally finite, there are finitely many edges e such that X_e intersects $\overline{B}_R^v(x)$. For $X_e \cap \overline{B}_R^v(x) \neq \emptyset$, the set $X_e \cap \overline{B}_R^v(x)$ is finite, since X_e is closed and discrete and $\overline{B}_R^v(x)$ is compact, so $\overline{B}_{R-1}^{\tau(e)}(\alpha_e(X_e \cap \overline{B}_R^v(x)))$ is compact.

Let $N_{1,1}, \dots, N_{1,k_1}$ be these finitely many compact sets. For each $N_{1,i}$ there are finitely many e such that X_e intersects $N_{1,i}$. Let $N_{1,i,1}, \dots, N_{1,i,k_{1,i}}$ be the finitely many closed $R-2$ balls about the sets $\alpha_e(X_e \cap N_{1,i})$ in their respective vertex spaces. Continue in this way for R steps. Let N be the union of the finitely many N_* , along with the finitely many rungs connecting them. This is a finite union of compact sets, so it is compact, and it contains $\overline{B}_R(x)$, since it contains every chain of length less than or equal to R starting from x . Thus, X is proper.

To see that X is geodesic, first consider two points x and x' contained in vertex spaces. By definition of the metric, there exists a sequence of chains joining x and x' with length decreasing to $d(x, x')$. Since x and x' are in vertex spaces, the length of a k -chain joining them is at least k , so there is a subsequence of approximating chains that are all k -chains for some fixed $k \leq d(x, x')$. Let $x_{0,i} := x, x'_{0,i}, \dots, x'_{k,i} := x'$ be the i -th k -chain in the sequence. Set $x_0 := x$. Now, $(x'_{0,i})$ is a bounded sequence in X_{v_0} . Since edge spaces are closed and discrete and edge patterns are locally finite, $(x'_{0,i})$ contains a constant subsequence. Pass to a subsequence of chains such that $(x'_{0,i})$ is constant and define x'_0 to be its value. Since edge patterns are locally finite, x'_0 belongs to finitely many X_e . Since attaching maps are coarsely invertible and edge spaces are closed and discrete, there are only finitely many rungs with x'_0 as an endpoint. Therefore, the sequence $(x_{1,i})$ takes only finitely many values. Choose one, define it to be x_1 , and pass to the subsequence with $x_{1,i} = x_1$. Repeat this process $k-1$ times, passing each time to a subsequence of chains with one additional constant coordinate. We are left with a constant sequence of chains with length decreasing to $d(x, x')$, so the length is equal to $d(x, x')$. Each vertex space and rung are geodesic spaces, so we connect successive points in the chain by geodesics in the vertex space or rung containing them to get a path with length equal to length of the chain, which is therefore a geodesic from x to x' .

If x' is an interior point of a rung and x is in a vertex space then find geodesics from x to the two endpoints of the rung and concatenate with the subsegment of the rung leading to x' . The shorter of these two paths is a geodesic. A similar argument works if both x and x' are interior points of rungs. \square

2.5 Algebraic trees of spaces

Let Γ be a graph of finitely generated groups. In this section we construct a tree of spaces over $T := T(\Gamma)$ that is quasi-isometric to $G := G(\Gamma)$. The idea is to take the vertex spaces to be Cayley graphs of the vertex stabilizers and use the edge injections of Γ to define the attaching maps. The construction is standard, but there is some bookkeeping involved that will be useful in Section 8.2.

For each $v \in \mathcal{V}\Gamma$, choose a finite generating set for G_v and coset representatives $h_{(v,i)}$ for G/G_v . For each $e \in \mathcal{E}\Gamma$ choose coset representatives $g_{(e,i)}$ of $G_{\iota(e)}/G_e$.

For each $v \in \mathcal{V}\Gamma$, choose a lift $\tilde{v} \in \mathcal{V}T$. For each edge $e \in \mathcal{E}\Gamma$, choose a lift $\tilde{e} \in \mathcal{E}T$ with $\iota(\tilde{e}) = \tilde{\iota(e)}$. Define $f_e := h_{(\iota(e),j)}g_{(e,i)}$ for i and j such that $f_e\tilde{e} = \tilde{e}$. Given a maximal subtree of Γ it is possible to choose lifts so that $f_e = 1$ for all edges e such that e or \bar{e} belongs to the maximal subtree.

For $t \in \mathcal{V}T \cup \mathcal{E}T$, let $\mathbf{t} \in \Gamma$ denote the image of t under the quotient by the G -action.

Let v be a vertex of T . There is a representative $h_{(v,i)}$ such that $v = h_{(v,i)}\tilde{v} \in T$. Define Y_v to be a copy of the Cayley graph of G_v with respect to the chosen generating set, which we identify with the coset $h_{(v,i)}G_v$ via left multiplication by $h_{(v,i)}$.

Take the edge spaces to be cosets of the edge stabilizers of Γ , and define attaching maps via containment. Specifically, for an edge $e = h_{(v,i)}g_{(e,j)}\tilde{e} \in T$ with $\iota(e) = v$ we define $Y_e := h_{(v,i)}g_{(e,j)}G_e \subset h_{(v,i)}G_v = Y_v$ with attaching map:

$$\alpha_e(x) := h_{(v,i)}g_{(e,j)}f_e\epsilon_e(g_{(e,j)}^{-1}h_{(v,i)}^{-1}x) \subset Y_{\tau(e)}$$

Let $Y := Y(T, \{Y_t\}, \{\alpha_e\})$ be the resulting tree of spaces over T , which we call an *algebraic tree of spaces* for Γ .

Lemma 2.15. *An algebraic tree of spaces Y for Γ is quasi-isometric to $G(\Gamma)$.*

Proof. By Lemma 2.14, Y is a proper geodesic metric space. The group $G(\Gamma)$ acts by left multiplication, properly discontinuously and cocompactly by isometries on Y , so they are quasi-isometric, by the Milnor-Švarc Lemma. \square

2.6 Trees of maps

2.6.1 Trees of quasi-isometries

Proposition 2.16. *Suppose X and X' are trees of spaces over T and T' , respectively. Suppose that $\chi: T \rightarrow T'$ is an isomorphism. Suppose there exists $M \geq 1$ and $A \geq 0$ such that:*

- *For each $v \in \mathcal{V}T$ there is an (M, A) -quasi-isometry $\phi_v: X_v \rightarrow X'_{\chi(v)}$ and a quasi-isometry inverse $\bar{\phi}_v: X'_{\chi(v)} \rightarrow X_v$.*
- *For every $e \in \mathcal{E}T$, the space $\phi_{\iota(e)}(X_e)$ is A -coarsely equivalent to $X'_{\chi(e)}$ in $X'_{\chi(\iota(e))}$, and the space $\bar{\phi}_{\iota(e)}(X'_{\chi(e)})$ is A -coarsely equivalent to X_e in $X_{\iota(e)}$.*

- For every $e \in \mathcal{ET}$ and every $x \in X_e$ there exists a point $x' \in X'_{\chi(e)}$ such that $d(\phi_{\iota(e)}(x), x') \leq A$ and $d(\phi_{\tau(e)}(\alpha_e(x)), \alpha'_e(x')) \leq A$.
- For every $e \in \mathcal{ET}$ and every $x' \in X'_{\chi(e)}$ there exists a point $x \in X_e$ such that $d(\bar{\phi}_{\iota(e)}(x'), x) \leq A$ and $d(\bar{\phi}_{\tau(e)}(\alpha'_e(x')), \alpha_e(x)) \leq A$.

Then there is a quasi-isometry $\phi: X \rightarrow X'$ with $\phi|_{X_v} = \phi_v$ for each vertex $v \in T$.

Definition 2.17. A collection of quasi-isometries (ϕ_v) satisfying the conditions given in Proposition 2.16 is called a *tree of quasi-isometries over χ compatible with X and X'* .

Proof of Proposition 2.16. It suffices to consider the unions of the vertex spaces, which form coarsely dense subsets of X and X' , and define ϕ by $\phi|_{X_v} := \phi_v$. Define a coarse inverse $\bar{\phi}$ to ϕ by $\bar{\phi}|_{X'_{\chi(v)}} := \bar{\phi}_v$.

Suppose x and x' are points in vertex spaces of X , and let $x_0, x'_0, x_1, \dots, x'_k$ be a k -chain joining them with length at most $d(x, x') + 1$. Let X_{v_i} be the vertex space containing x_i and x'_i , and let e_i be the edge such that x'_{i-1} and x_i are endpoints of a rung over e_i .

Let $y_0 := \phi_{v_0}(x_0)$ and $y'_k := \phi_{v_k}(x'_k)$. By hypothesis, for each x'_i there exists a $y'_i \in X'_{\chi(e_i)}$ such that $d(\phi_{\iota(e)}(x'_i), y'_i) \leq A$ and $d(\phi_{\tau(e)}(x_{i+1}), y_{i+1}) \leq A$, where $y_{i+1} := \alpha'_{\chi(e)}(y'_i)$. This gives a k -chain y_0, y'_0, \dots, y'_k joining $\phi(x)$ and $\phi(x')$ whose length is:

$$\begin{aligned}
& d_{X'_{\chi(v_0)}}(y_0, y'_0) + \sum_{0 < i \leq k} (d(y'_{i-1}, y_i) + d_{X'_{\chi(v_i)}}(y_i, y'_i)) \\
& \leq d_{X'_{\chi(v_0)}}(\phi_{v_0}(x_0), \phi_{v_0}(x'_0)) + A + \\
& \quad \sum_{0 < i \leq k} (1 + 2A + d_{X'_{\chi(v_i)}}(\phi_{v_i}(x_i), \phi_{v_i}(x'_i))) \\
& \leq M d_{X_{v_0}}(x_0, x'_0) + 2A + \sum_{0 < i \leq k} (1 + 3A + M d_{X_{v_i}}(x_i, x'_i)) \\
& = 2A + k(1 + 3A) + M \sum_{0 \leq i \leq k} d_{X_{v_i}}(x_i, x'_i) \\
& \leq 2A + \max\{1 + 3A, M\} \left(k + \sum_{0 \leq i \leq k} d_{X_{v_i}}(x_i, x'_i) \right) \\
& \leq 2A + \\
& \quad \max\{1 + 3A, M\} \left(d_{X_{v_0}}(x_0, x'_0) + \sum_{0 < i \leq k} (d(x'_{i-1}, x_i) + d_{X_{v_i}}(x_i, x'_i)) \right) \\
& \leq 2A + \max\{1 + 3A, M\} (d(x, x') + 1)
\end{aligned}$$

Therefore, $d(\phi(x), \phi(x')) \leq \max\{1 + 3A, M\} d(x, x') + (2A + \max\{1 + 3A, M\})$, so ϕ is coarsely Lipschitz. A similar computation shows $\bar{\phi}$ is coarsely Lipschitz, so ϕ is a quasi-isometry. \square




Figure 1: Commuting diagram for Corollary 2.18.

Corollary 2.18. *Suppose $\chi: T \rightarrow T'$ is an isomorphism of trees and X and X' are trees of spaces over T and T' , respectively. Suppose there are $M \geq 1$ and $A \geq 0$ and (M, A) -quasi-isometries $\phi_v: X_v \rightarrow X'_{\chi(v)}$ for each vertex and $\phi_e: X_e \rightarrow X'_{\chi(e)}$ for each edge such that the diagram in Figure 1 commutes up to uniformly bounded error. Then (ϕ_v) is a tree of quasi-isometries over χ compatible with X and X' , so X is quasi-isometric to X' .*

2.6.2 Gromov boundaries and trees of homeomorphisms

In this section, let X be a proper geodesic hyperbolic tree of spaces over a tree T , such that the vertex spaces are proper geodesic hyperbolic spaces, the edge spaces are uniformly quasi-convex in X , and the attaching maps are uniform quasi-isometries.

For example, if Γ is a finite acylindrical graph of hyperbolic groups such that the edge injections are quasi-isometric embeddings then the algebraic tree of spaces has this structure [4, 24]. If, moreover, the edge groups are two-ended then the edge injections are automatically quasi-isometric embeddings.

Quasi-convexity of edge spaces implies quasi-convexity of vertex spaces, so the Gromov boundary of each vertex space and edge space embeds into the boundary of X .

Consider the space $\partial X := \partial T \sqcup \coprod_{v \in \mathcal{V}T} \partial X_v$ modulo identifying $\partial X_e \subset \partial X_{\iota(e)}$ with $\partial X_{\bar{e}} \subset \partial X_{\tau(e)}$ via $\partial \psi_e$ for each edge e . Let ∂_{stab} denote the image of $\coprod_{v \in \mathcal{V}T} \partial X_v$ in ∂X .

Lemma 2.19. *The inclusion of ∂X into the Gromov boundary of X is a surjection.*

Proof. Pick a basepoint $p \in X$ and let $\zeta: [0, \infty] \rightarrow \bar{X}$ be a geodesic ray based at p . Consider the projection $\pi(p)$ of p to T .

If $\pi \circ \zeta$ crosses some edge e infinitely many times then there is an unbounded sequence of times t_1, t_2, \dots such that $\zeta(t_i) \in X_e$. Since the edge space is quasi-convex, ζ stays bounded distance from X_e , so converges to a point in ∂X_e .

A similar argument shows that if $\pi \circ \zeta$ visits a vertex v infinitely many times or is eventually constant at v then ζ converges to a point in ∂X_v .

The remaining possibility is that $\pi \circ \zeta$ limits to a point of ∂T . We claim there is a unique asymptotic class of geodesic rays in X with $\pi \circ \zeta \rightarrow \eta \in \partial T$. Suppose ζ' is another

such ray. Let v_0 be the vertex such that $p \in X_{v_0}$. Let e be an edge in T on the geodesic from v_0 to η . Let e' be another edge on the geodesic from v_0 to η that is distance at least $\delta + Q$ from e , where δ is the hyperbolicity constant for X and Q is the quasi-convexity constant for edge spaces. Let x be a point of $\zeta \cap X_e$ and y a point of $\zeta \cap X_{e'}$. Let x' be a point of $\zeta' \cap X_e$ and y' a point of $\zeta' \cap X_{e'}$. Consider a geodesic triangle whose sides are $\zeta|_{[p,y]}$, $\zeta'|_{[p,y']}$, and a geodesic ζ'' from y to y' . By quasi-convexity, $d(x, \zeta'') > \delta$ and $d(x', \zeta'') > \delta$. Therefore, hyperbolicity implies $d(x, \zeta') \leq \delta$ and $d(x', \zeta) \leq \delta$. This implies ζ and ζ' are asymptotic. \square

Now we define a topology on ∂X and show it is equivalent to the Gromov topology.

Definition 2.20 (domains). The *domain* $D(\eta)$ of a point $\eta \in \partial T$ is the singleton $\{\eta\}$. The *domain* $D(\xi)$ of a point ξ of $\partial_{Stab} X$ is the subtree of T spanned by those vertices v of T such that ∂X_v contains a point in the equivalence class ξ .

Remark 2.21. It can be proved that domains of points of $\partial_{Stab} X$ have a uniformly bounded number of edges, see [30].

We can now define neighborhoods of points of ∂X , starting with points of ∂T .

Definition 2.22 (neighborhoods of points of ∂T). Let η be a point of ∂T , and U be a neighborhood of η in $T \cup \partial T$. We define the neighborhood $V_U(\eta)$ as the set of points of ∂X whose domain is contained in U .

Before moving to neighborhoods of points of $\partial_{Stab} X$, we need a definition.

Definition 2.23 (ξ -family, cone). Let ξ be a point of $\partial_{Stab} X$. For every vertex v of $D(\xi)$, choose a neighborhood U_v of ξ in $X_v \cup \partial X_v$. Let \mathcal{U} be the collection of sets $U_v, v \in D(\xi)$, which we call a ξ -family. We define the set $\text{Cone}_{\mathcal{U}}(\xi)$, called a *cone*, as the set of points $w \in (T \cup \partial T) \setminus D(\xi)$ such that if e is the last edge of the geodesic from w to $D(\xi)$ in T , we have $\partial \alpha_e(\partial X_e) \subset U_{\tau(e)}$.

Definition 2.24 (neighborhoods of points of $\partial_{Stab} X$). Let ξ be a point of $\partial_{Stab} X$, and \mathcal{U} be a ξ -family. We define the neighborhood $V_{\mathcal{U}}(\xi)$ as the set of points η of ∂X such that the following holds:

- $D(\eta) \setminus D(\xi)$ is contained in $\text{Cone}_{\mathcal{U}}(\xi)$,
- for every vertex v of $D(\xi) \cap D(\eta)$, we have $\eta \in U_v$.

Theorem 2.25 ([30, Corollary 9.19]). *With the topology described above, the inclusion of ∂X into the Gromov boundary of X is a homeomorphism.* \square

Definition 2.26. Let X and X' be proper geodesic hyperbolic trees of quasi-convex spaces over a trees T and T' , respectively. A *tree of boundary homeomorphisms compatible with X and X' over an isomorphism $\chi: T \rightarrow T'$* consists of homeomorphisms $\rho_v: \partial X_v \rightarrow \partial X'_{\chi(v)}$ for every vertex $v \in T$ such that for $\xi \in \partial X_v \cap \partial X_w$ we have $\rho_v(\xi) = \rho_w(\xi) \in \partial X'_{\chi(v)} \cap \partial X'_{\chi(w)}$, and for $\xi \in \partial X'_v \cap \partial X'_w$ we have $\rho_v^{-1}(\xi) = \rho_w^{-1}(\xi) \in \partial X_{\chi^{-1}(v)} \cap \partial X_{\chi^{-1}(w)}$.

Proposition 2.27. *Let X and X' be proper geodesic hyperbolic trees of quasi-convex spaces over trees T and T' , respectively, with a compatible tree of boundary homeomorphisms (ρ_v) over $\chi \in \text{Isom}(T, T')$. Then there is a homeomorphism $\rho: \partial X \rightarrow \partial X'$ defined by $\rho|_{\partial X_v} := \rho_v$ and $\rho|_{\partial T} := \partial\chi$.*

Proof. We have a well defined map ρ and its inverse ρ^{-1} is defined by $\rho^{-1}|_{\partial X'_v} := \rho_{\chi^{-1}(v)}^{-1}$ and $\rho^{-1}|_{\partial T} := \partial\chi^{-1}$. It is clear that ρ and ρ^{-1} are continuous with respect to the topology given in Definition 2.22 and Definition 2.24, which is equivalent to the standard topology by Theorem 2.25. \square

Theorem 2.28. *Let G and G' be one-ended hyperbolic groups with non-trivial two-ended JSJ decompositions. Let X and X' be algebraic trees of spaces over the respective JSJ trees of cylinders $T := \text{Cyl}(G)$ and $T' := \text{Cyl}(G')$. Every homeomorphism $\rho: \partial X \rightarrow \partial X'$ splits as a tree of compatible boundary homeomorphisms over the isomorphism $\rho_*: T \rightarrow T'$. Every tree of boundary homeomorphisms (ρ_v) compatible with X and X' over an isomorphism $\chi: T \rightarrow T'$ patches together to give a homeomorphism $\rho: \partial X \rightarrow \partial X'$ with $\rho_* = \chi$ and $\rho|_{\partial X_v} = \rho_v$.*

Proof. Since the edge groups are two-ended they are virtually cyclic, hence quasi-convex. Therefore, the boundary of each vertex space embeds into the boundary of its tree of spaces. By Theorem 2.11, ρ induces an isomorphism $\rho_*: T \rightarrow T'$ and $\rho(\partial X_v) = \partial X'_{\rho_*(v)}$. Since these spaces are embedded, $\rho|_{\partial X_v} \in \text{Homeo}((\partial X_v, \partial\mathcal{P}_v), (\partial X'_{\rho_*(v)}, \partial\mathcal{P}'_{\rho_*(v)}))$, so ρ splits as a tree of boundary homeomorphisms over ρ_* .

The converse is Proposition 2.27. \square

3 Distinguishing orbits: Basic examples

Given a group G , a common strategy in understanding groups quasi-isometric to G is to first understand self-quasi-isometries of G . We explained in the previous section that $\mathcal{QI}(G)$ acts on $\text{Cyl}(G)$, preserving the relative quasi-isometry types of vertices. A first step towards understanding $\mathcal{QI}(G)$ is to understand its action on $\text{Cyl}(G)$. In particular, we would like to determine the $\mathcal{QI}(G)$ -orbits of vertices in $\text{Cyl}(G)$. We know that there are finitely many G -orbits of vertices in $\text{Cyl}(G)$, and that $\mathcal{QI}(G)$ -orbits are unions of G -orbits, so the problem reduces to distinguishing G -orbits that are not contained in a common $\mathcal{QI}(G)$ -orbit.

In this section we show on some particular examples how one may distinguish G -orbits. These examples motivate the technical machinery of the following sections, in which we combine and iterate the considerations presented here.

Our examples are hyperbolic, so, as usual, we consider the cases of quasi-isometry and boundary homeomorphism simultaneously by letting $\text{Map}(G)$ mean either $\mathcal{QI}(G)$ or $\text{Homeo}(\partial G)$, as appropriate.

In general our techniques reduce questions about graphs of groups to questions about the vertex groups, so our results are strongest, and the examples are most illuminating,

when the vertex groups are well understood. First we establish such a well understood vertex group, and then we use it to construct graph of group examples.

Consider the free group $F_2 = \langle a, b \rangle$. Let X be the Cayley graph of F_2 with respect to the generating set $\{a, b\}$. Recall that if \mathcal{H} is a finite collection of two-ended subgroups of F_2 then there is a corresponding peripheral structure \mathcal{P} on X consisting of distinct coarse equivalence classes of conjugates of elements of \mathcal{H} . Since X is a tree, the coarse equivalence class of a two-ended subgroup contains a unique bi-infinite geodesic, so we can think of \mathcal{P} as an F_2 -invariant collection of bi-infinite geodesics.

We claim that $\llbracket (X, \mathcal{P}_i) \rrbracket = \llbracket (X, \mathcal{P}_j) \rrbracket$ for all i and j where \mathcal{P}_i is the peripheral structure induced by \mathcal{H}_i defined as follows:

- $\mathcal{H}_0 = \{\langle a^2 b^2 \bar{a} \bar{b} \rangle\}$
- $\mathcal{H}_1 = \{\langle ab^2 \rangle, \langle a^2 \bar{b} \rangle\}$
- $\mathcal{H}_2 = \{\langle ab \rangle, \langle a^2 \bar{b}^2 \rangle\}$
- $\mathcal{H}_3 = \{\langle a \rangle, \langle ab \bar{a} \bar{b}^2 \rangle\}$
- $\mathcal{H}_4 = \{\langle a \rangle, \langle b \rangle, \langle ab \bar{a} \bar{b} \rangle\}$
- $\mathcal{H}_5 = \{\langle a \rangle, \langle b \rangle, \langle ab \rangle, \langle a \bar{b} \rangle\}$

This follows from [14, Theorem 6.3]. The reason is that the Whitehead graph for each of these examples is the complete graph on four vertices. For each i , there are exactly 6 bi-infinite geodesics representing elements of \mathcal{P}_i passing through the vertex of X corresponding to the identity element of F_2 . Given i and j , any choice of bijection between these six elements of \mathcal{P}_i and the six corresponding elements of \mathcal{P}_j extends uniquely to an element of $\mathcal{QI}((X, \mathcal{P}_i), (X, \mathcal{P}_j))$. Conversely, every element of $\mathcal{QI}((X, \mathcal{P}_i), (X, \mathcal{P}_j))$ is of this form, up to pre- and post-composition with the F_2 -action.

When F_2 is a vertex group of a graph of groups and the peripheral structure coming from incident edge groups is one of the \mathcal{P}_i above then the vertex is rigid.

For variety, we give two more rigid examples.

- $F_3 = \langle a, b, c \rangle$ relative to the peripheral structure induced by $\langle a^2 b^2 c^2 a c b \rangle$. The Whitehead graph for this example is the complete bipartite graph with parts of three vertices each. This peripheral structure is not quasi-isometrically equivalent to the \mathcal{P}_i given above.
- $\pi_1(M)$, for M a closed hyperbolic 3-manifold, relative to the peripheral structure induced by $\langle g \rangle$, where g is a non-trivial, indivisible element of $\pi_1(M)$.

3.1 First example: wrong type, bad neighbors, stretch factors

Consider the graph of groups $\mathbf{\Gamma}$ shown in Figure 2. It is the graph of cylinders for the one-ended hyperbolic group $G := G(\mathbf{\Gamma})$, so $T := T(\mathbf{\Gamma}) = \text{Cyl}(G)$. The label below a vertex is its name, and the label above a vertex is the associated vertex group. All of the edge groups are infinite cyclic. Suppose that we have chosen a generator for each infinite cyclic vertex and edge group. Each end of an edge is labelled to indicate the image of this generator in the incident vertex group. For cyclic vertices we give a non-zero integer indicating that the image of the edge generator is that power of the vertex generator.




Figure 2: First example Γ

This label is omitted when it is equal to 1, which is the case for all edges in this particular graph of groups.

For non-cyclic vertices we indicate the element that is the image of the chosen generator of the edge group. For economy of notation we reuse a and b as generators of all the F_2 vertex groups. These are distinct subgroups of G , so more properly we would say, for instance, that $G_{r_1} = \langle a_1, b_1 \rangle$ and the two incident edge injections send the generators of their respective edge groups to $a_1 b_1^2$ and $a_1 \bar{b}_1^2$.

In this section we describe the $\text{Map}(G)$ -orbits of vertices in T .

3.1.1 Wrong type

The vertex h is the only hanging vertex, so $G\tilde{h}$ is a $\text{Map}(G)$ -orbit.

The vertices labelled c_i are cylindrical, while the vertices labelled r_i are rigid, so there are no i and j such that \tilde{r}_i and \tilde{c}_j are in the same $\text{Map}(G)$ -orbit.

The vertex group of r_6 is not quasi-isometric or boundary homeomorphic to any of the other vertex groups, so $G\tilde{r}_6$ is a $\text{Map}(G)$ -orbit.

The vertex group of r_9 relative to the peripheral structure induced by the incident edge is not the same relative quasi-isometry or boundary homeomorphism type of any other vertex, so $G\tilde{r}_9$ is a $\text{Map}(G)$ -orbit.

3.1.2 Bad neighbors

The cylindrical vertices have finite valence in T . The valence of vertices in $G\tilde{c}_1$ is 5, while for all other cylindrical vertices it is 2. Thus, $G\tilde{c}_1$ is a $\text{Map}(G)$ -orbit.

Vertices in $G\tilde{c}_2$ are the only ones adjacent to vertices in the $\text{Map}(G)$ -orbit $G\tilde{r}_6$, so $G\tilde{c}_2$ is a $\text{Map}(G)$ -orbit.

Vertices in $G\tilde{c}_2$ are adjacent only to vertices in $G\tilde{r}_1$ and $G\tilde{r}_6$, but we already know that $G\tilde{r}_6$ is a $\text{Map}(G)$ -orbit, so $G\tilde{r}_1$ is a $\text{Map}(G)$ -orbit as well.

By the same sort of argument, we quickly conclude that $G\tilde{c}_5$, $G\tilde{r}_4$, $G\tilde{c}_6$, and $G\tilde{r}_5$ are distinct $\text{Map}(G)$ -orbits.

Such arguments based on distinguishing vertices by their neighbors are the contents of Section 4.

3.1.3 Stretch factors

It remains to differentiate the second and third branches of $\mathbf{\Gamma}$. This is more subtle, as the only difference is at vertices r_2 and r_3 , which do have the same relative quasi-isometry type.

Notices that the generator of G_{c_3} is identified with an element of word length 6 in G_{r_7} , and length 3 in G_{r_2} , whereas the generator of G_{c_4} is identified with elements of word lengths 6 and 2. It is not at all obvious that these lengths are preserved by quasi-isometries. Nevertheless, we will show that in this particular example \tilde{c}_3 and \tilde{c}_4 can be distinguished up to quasi-isometry by the fact that the ratios $\frac{6}{3}$ and $\frac{6}{2}$ are not equal. These ratios correspond to an invariant called the ‘stretch factor’, which is developed in Section 5.

By neighbor arguments as above we can then conclude that $\mathcal{QI}(G)$ -orbits of T are the same as G -orbits.

The stretch factor is only an invariant for quasi-isometries, not boundary homeomorphism. In this example we cannot, in fact, distinguish the second and third branch of $\mathbf{\Gamma}$ up to boundary homeomorphism, which is to say, $G\tilde{r}_2 \cup G\tilde{r}_3$ is a single $\text{Homeo}(\partial G)$ -orbit, as are $G\tilde{c}_3 \cup G\tilde{c}_4$ and $G\tilde{r}_7 \cup G\tilde{r}_8$.

3.2 Second example: unbalanced cylinders

In general, if \mathcal{L} is an element of a peripheral structure \mathcal{P} on X , there is no reason to believe that there is an element of $\text{Map}((X, \mathcal{P}))$ that preserves \mathcal{L} and reverses its orientation.

We briefly give an idea how one may construct such an example. Consider the genus 7 surface Σ in Figure 3, with a hyperbolic metric.

Let $f \in \pi_1(\Sigma)$ be the element represented by the curve running around the central hole. Let \mathcal{L} be a component of the preimage of this curve in $\tilde{\Sigma} = \mathbb{H}^2$. The figure shows three curves that separate Σ into subsurfaces Σ_1 , Σ_2 , and Σ_3 , each of which is a genus two surface with two boundary components. Suppose each Σ_i contains a collection of curves that is complicated enough so that all the complementary components are discs, and different enough so that there does not exist a crossing-preserving bijection between the components of the preimages of the curves of Σ_i contained in one component of the preimage of Σ_i in $\tilde{\Sigma}$ and the components of the preimages of the curves of Σ_j contained in one component of the preimage of Σ_j . Let \mathcal{P} be the peripheral structure on $\tilde{\Sigma}$ induced by all of the above curves in Σ , which is to say, each element of \mathcal{P} is the coarse equivalence class of a component of the preimage in $\tilde{\Sigma}$ of one of the curves in Σ .




Figure 3: The surface Σ

Then there is no element of $\text{Map}(\pi_1(\Sigma), \mathcal{P})$ that reverses \mathcal{L} . The reason is that \mathcal{L} passes through components of the preimages of Σ_1 , Σ_2 , and Σ_3 in order. If this order were reversible by an element of $\text{Map}(\pi_1(\Sigma), \mathcal{P})$ we would contradict the restriction that the pattern of curves in Σ_1 , Σ_2 , and Σ_3 were chosen to be ‘different’.

Now consider the graph of groups $\mathbf{\Gamma}$ of Figure 4, where Σ and f are as above, and g is a non-trivial, indecomposable element of a closed hyperbolic 3-manifold M . Then $G := G(\mathbf{\Gamma})$ is a one-ended hyperbolic group with $T := T(\mathbf{\Gamma}) = \text{Cyl}(G)$.




Figure 4: Second example $\mathbf{\Gamma}$

The graph of groups $\mathbf{\Gamma}$ has a central cylindrical vertex c_1 that attaches to three rigid vertices r_1 , r_2 , and r_3 , each of which has local group $\pi_1(\Sigma)$, with each attachment along a copy of f . For each r_i there are some number of other incident edges, each corresponding to one of the curves in Σ described above, so that the peripheral structure on G_{r_i} induced by edge inclusions is the peripheral structure \mathcal{P} described above for $\tilde{\Sigma}$. Each of these edges we attach to another cylindrical vertex, and then to a rigid vertex carrying a copy of $\pi_1(M)$.

We consider a kind of ‘ f -parity’ by counting the number of vertices adjacent to \tilde{c}_1 for which the generator of $G_{\tilde{c}_1}$ is identified with f minus the number of vertices adjacent to \tilde{c}_1 for which the generator of $G_{\tilde{c}_1}$ is identified with \bar{f} . Since \tilde{c}_1 has non-zero f -parity, we call it an *unbalanced cylinder*. We claim that there is not an element of $\text{Map}(G)$ that reverses the orientation of an unbalanced cylinder. In this example, we then conclude that $G\tilde{r}_3$ is a $\text{Map}(G)$ -orbit and $G\tilde{r}_1 \cup G\tilde{r}_2$ is a different $\text{Map}(G)$ -orbit.

The proof is as follows. Since $G\tilde{c}_1$ is the only G -orbit of cylindrical vertices in T of valence 3, if there were an element of $\text{Map}(G)$ taking \tilde{r}_3 to, say, \tilde{r}_1 , then it would fix \tilde{c}_1 . There is clearly an element of $\text{Map}((G_{r_3}, \mathcal{P}_{r_3}), (G_{r_1}, \mathcal{P}_{r_1}))$ taking f to \bar{f} , but, since f is not reversible, if such a map extended to an element of $\text{Map}(G)$ it would necessarily reverse the orientation of $G_{\tilde{c}_1}$. Since f is not reversible, this would mean that every vertex in which f is identified with the generator of $G_{\tilde{c}_1}$ must be sent to a vertex in which \bar{f} is identified with the generator of $G_{\tilde{c}_1}$. This means that both \tilde{r}_1 and \tilde{r}_2 are sent to \tilde{r}_3 , contradicting the fact that $\text{Map}(G)$ acts by isomorphisms on T .

Dealing with issues of orientation and unbalanced cylinders is the focus of Section 6.2.

4 Decorated trees and structure invariants

Let G be a group and let T be a simplicial tree of countable valence upon which G acts cocompactly and without inverting an edge. In this section, we explain how to associate to T an invariant, the *structure invariant*, that completely characterizes it up to decoration-preserving isomorphism. The construction of this invariant generalizes the ‘degree refinement’ invariant of graph theory, which is an invariant that determines when two finite graphs have isomorphic universal covers.

Definition 4.1. A *decoration* is a G -invariant map $\delta: T \rightarrow \mathcal{O}$ that assigns to each vertex of T an *ornament* $o \in \mathcal{O}$.

For simplicity, in this section we decorate only vertices. In the next section we will also decorate edges, with the condition that $\delta(e) = \delta(\bar{e})$ for every edge. Formally, this can be accomplished by subdividing each edge of T and decorating the new vertices.

Corresponding to a decoration there is a partition of T as $\coprod_{o \in \mathcal{O}} \delta^{-1}(o)$. We say that a decoration $\delta': T \rightarrow \mathcal{O}'$ is a *refinement* of δ if the δ' -partition is finer than the δ -partition. Equivalently, a decoration $\delta': T \rightarrow \mathcal{O}'$ is a refinement of $\delta: T \rightarrow \mathcal{O}$ if there exists a surjective map $\pi: \text{Im } \delta' \rightarrow \text{Im } \delta$ such that $\pi \circ \delta' = \delta$. We say δ' is a *strict refinement* if the δ' -partition is strictly finer than the δ -partition. A refinement that is not strict is a *trivial refinement*.

In this section, we will define a *structure invariant* for a decorated tree, such that we can reconstruct the tree, up to decoration-preserving tree automorphism, from the data of the structure invariant. In order to do so, we want to identify orbits in T under the action of the group $\text{Aut}(T, \delta) := \{\chi \in \text{Aut}(T) \mid \delta \circ \chi = \delta\}$, and then to say how they fit together. In a nutshell, the idea is as follows.

Suppose v and w are two vertices. If there is a $\chi \in \text{Aut}(T, \delta)$ with $\chi(v) = w$, then, of course, $\delta(v) = \delta(w)$. Additionally, χ gives a decoration-preserving bijection from the

neighbors of v to the neighbors of w . Thus, for each ornament o , the number of neighbors of v bearing o must be equal to the number of neighbors of w bearing o .

Conversely, if $\delta(v) = \delta(w)$, but for some ornament o there are differing numbers of neighbors of v bearing o and neighbors of w bearing o , then there is no decoration-preserving automorphism taking v to w , so we ought refine the decoration to distinguish v from w . We then repeat this refinement process until vertices with the same ornament can no longer be distinguished by the ornaments of their neighbors. This happens after finitely many steps because $G \backslash T$ is compact, see Proposition 4.3. Section 4.1 formalizes this process, which we call *neighbor refinement*.

4.1 Neighbor refinement

Let $\overline{\mathbb{N}} := \mathbb{N} \cup \{0, \infty\}$. Call $\mathcal{O}_0 := \mathcal{O}$ and $\delta_0 := \delta$ the ‘initial set of ornaments’ and the ‘initial decoration’, respectively. Beginning with $i = 0$, for each $v \in \mathcal{VT}$ define:

$$f_{v,i}: \mathcal{O}_i \rightarrow \overline{\mathbb{N}}: o \mapsto \#\{w \in \delta_i^{-1}(o) \mid w \text{ is adjacent to } v\}$$

Define $\mathcal{O}_{i+1} := \mathcal{O}_0 \times \overline{\mathbb{N}}^{\mathcal{O}_i}$, and define $\delta_{i+1}(v) := (\delta_0(v), f_{v,i})$.

Lemma 4.2. *For all i , the map $\delta_{i+1}: T \rightarrow \mathcal{O}_{i+1}$ is a decoration that refines $\delta_i: T \rightarrow \mathcal{O}_i$.*

Proof. Let v be a vertex, and let $g \in G$. Suppose δ_i is G -invariant. Then, $\delta_{i+1}(gv) = (\delta_i(gv), f_{gv,i}) = (\delta_i(v), f_{v,i}) = \delta_{i+1}(v)$. Since δ_0 is G -invariant, all the δ_i are decorations by induction.

For each i , let $N(v, i, o)$ denote the number of neighbors of v in $\delta_i^{-1}(o)$.

Clearly, δ_1 refines δ_0 , since δ_0 is the composition of δ_1 with projection to the first coordinate of the image. Suppose that δ_i refines δ_{i-1} . Then for every $o' \in \mathcal{O}_{i-1}$, we have $N(v, i-1, o') = \sum_{o \in \delta_i \circ \delta_{i-1}^{-1}(o')} N(v, i, o)$.

If $\delta_{i+1}(v) = \delta_{i+1}(w)$ then $\delta_0(v) = \delta_0(w)$ and $N(v, i, o) = N(w, i, o)$ for each $o \in \mathcal{O}_i$. Thus, for all $o' \in \mathcal{O}_{i-1}$,

$$N(v, i-1, o') = \sum_{o \in \delta_i \circ \delta_{i-1}^{-1}(o')} N(v, i, o) = \sum_{o \in \delta_i \circ \delta_{i-1}^{-1}(o')} N(w, i, o) = N(w, i-1, o'),$$

so $\delta_i(v) = \delta_i(w)$. Hence δ_{i+1} refines δ_i . The lemma follows by induction. \square

Proposition 4.3. *There exists an $s \geq 0$ such that δ_{i+1} is a strict refinement of δ_i for all $i+1 \leq s$ and δ_{i+1} is a trivial refinement of δ_i for all $i \geq s$.*

Proof. Since δ_{i+1} refines δ_i there is a function $\delta_i \circ \delta_{i+1}^{-1}: \text{Im } \delta_{i+1} \subset \mathcal{O}_{i+1} \rightarrow \mathcal{O}_i$. The refinement is trivial precisely when this function is injective.

Suppose there exists an s such that δ_{s+1} is a trivial refinement of δ_s . For every $o \in \text{Im } \delta_{s+1} \subset \mathcal{O}_{s+1}$ we have $f_{v,s+1}(o) = f_{v,s} \circ \delta_s \circ \delta_{s+1}^{-1}(o)$, and $f_{v,s+1}(o) = 0$ otherwise. Therefore, $\delta_{s+1}(v) = \delta_{s+1}(w)$ implies $f_{v,s} = f_{w,s}$, which implies $f_{v,s+1} = f_{w,s+1}$, which implies $\delta_{s+2}(v) = \delta_{s+2}(w)$. Thus, once one refinement in the sequence is trivial so are all further refinements.

To see that eventually some refinement is trivial, note that G -invariance implies that for all i we have a partition of T by $T = \coprod_{o \in \mathcal{O}_i} \delta_i^{-1}(o)$ in which each part is a union of G -orbits. A refinement is strict if and only if the new partition has strictly more parts than the previous one. However, the number of parts is bounded above by the number of G -orbits, which is finite. \square

Definition 4.4. The *neighbor refinement* of δ is the decoration $\delta_s: T \rightarrow \mathcal{O}_s$ at which the neighbor refinement process stabilizes.

Proposition 4.5. Let $\delta': T \rightarrow \mathcal{O}'$ be the neighbor refinement of $\delta: T \rightarrow \mathcal{O}$. The δ' -partition of T is equal to the partition into $\text{Aut}(T, \delta)$ -orbits and to the partition into $\text{Aut}(T, \delta')$ -orbits.

Proof. The partition into $\text{Aut}(T, \delta)$ -orbits is finer than the δ' -partition because each refinement step is $\text{Aut}(T, \delta)$ -equivariant. The $\text{Aut}(T, \delta')$ -orbit partition is finer than the $\text{Aut}(T, \delta)$ -orbit partition since δ' is a refinement of δ . We show the δ' -partition is finer than the $\text{Aut}(T, \delta')$ -orbit partition by supposing $\delta'(v) = \delta'(w)$ and producing $\chi \in \text{Aut}(T, \delta')$ with $\chi(v) = w$.

The automorphism χ is constructed inductively. Start by defining $\chi(v) := w$. Since δ' is stable under neighbor refinement, for every $o \in \mathcal{O}'$ we have $\#\text{lk}(v) \cap (\delta')^{-1}(o) = \#\text{lk}(w) \cap (\delta')^{-1}(o)$. Extend χ by choosing any bijection between these sets. This extends χ to the 1-neighborhood of v .

Now suppose χ is defined on a subtree T' of T such that for every $v' \in \mathcal{V}T'$, either v' is a leaf or T' contains every edge incident to v' . Let v' be a leaf, and let u be the vertex of T' adjacent to v' . For $\delta'(u)$ extend χ via a bijection between $(\text{lk}(v) \setminus \{u\}) \cap (\delta')^{-1}(\delta'(u))$ and $(\text{lk}(\chi(v')) \setminus \{\chi(u)\}) \cap (\delta')^{-1}(\delta'(u))$. For $o \neq \delta'(u)$ extend χ just as in the base case. \square

4.2 Structure invariants

Let $\delta_s: T \rightarrow \mathcal{O}_s$ be a neighbor refinement of δ as in Proposition 4.3. As defined, δ_s actually encodes the history of the refinement process, not just the structure of T . We define the structure invariant by forgetting this extraneous history. Let $\pi_0: \mathcal{O}_s \rightarrow \mathcal{O}$ be projection to the first coordinate. Choose an ordering of $\text{Im } \delta$ and let $\mathcal{O}[j]$ denote the j -th ornament in the image. Similarly, for each $1 \leq j \leq \#\text{Im } \delta$, choose an ordering of $\pi_0^{-1}(\mathcal{O}[j]) \cap \text{Im } \delta_s$. Order $\text{Im } \delta_s$ lexicographically, and let $\mathcal{O}_s[i]$ denote the i -th ornament.

Definition 4.6. The *structure invariant* $\mathcal{S}(T, \delta, \mathcal{O})$ is the $\#\text{Im } \delta_s \times \#\text{Im } \delta_s$ matrix whose j, k -entry is the tuple consisting of the number of vertices in $\delta_s^{-1}(\mathcal{O}_s[j])$ adjacent to each vertex of $\delta_s^{-1}(\mathcal{O}_s[k])$, the \mathcal{O} -ornament $\pi_0(\mathcal{O}_s[j])$ and the \mathcal{O} -ornament $\pi_0(\mathcal{O}_s[k])$. The last two coordinates are called respectively the *row* and *column ornament* of the entry.

$\mathcal{S}(T, \delta, \mathcal{O})$ can be seen as a block matrix, with blocks consisting of entries with the same row ornaments and the same column ornaments. The structure invariant is well defined up to permuting the \mathcal{O} -blocks and permuting rows and columns within \mathcal{O} -blocks, ie, up to the choice of orderings of \mathcal{O} and the $\pi_0^{-1}(\mathcal{O}[j])$.

Proposition 4.7. *Let $\delta: T \rightarrow \mathcal{O}$ be a G -invariant decoration of a cocompact G -tree. Let $\delta': T' \rightarrow \mathcal{O}$ be a G' -invariant decoration of a cocompact G' -tree. There exists a decoration-preserving isomorphism $\phi: T \rightarrow T'$ if and only if $\mathcal{S}(T, \delta, \mathcal{O}) = \mathcal{S}(T', \delta', \mathcal{O})$, up to permuting rows and columns within \mathcal{O} -blocks.*

In particular, with the above notations, T and T' must have the same sets of ornaments for a decoration-preserving isomorphism ϕ to exist.

Proof. It is clear that isomorphic decorated trees have the same structure invariants, up to choosing the orderings of the ornaments. For the converse, assume that we have reordered within \mathcal{O} -blocks so that $\mathcal{S}(T, \delta, \mathcal{O}) = \mathcal{S}(T', \delta', \mathcal{O}) = \mathcal{S}$. Construct a decoration-preserving tree isomorphism exactly as in the proof of Proposition 4.5. \square

Remark. When T is the universal cover of a finite graph Γ and the initial set of ornaments is trivial then the structure invariant we have defined is just the well known *degree refinement* of Γ . The lemma says that two graphs have the same degree refinement if and only if they have isomorphic universal covers. A theorem of Leighton [27] says that such graphs in fact have a common finite cover. There are also decorated versions of Leighton's Theorem, eg [34].

Observation. We get a quasi-isometry invariant of a group G by taking the structure invariant of a cocompact $\mathcal{QI}(G)$ -tree with a $\mathcal{QI}(G)$ -invariant decoration.

This is a simple observation, but it does not seem to have appeared in the literature in this generality.

Behrstock and Neumann [2, 3] have used special cases of this type of invariant, in a different guise, to classify fundamental groups of some families of compact irreducible 3-manifolds of zero Euler characteristic. In both papers the tree is the Bass-Serre tree for the geometric decomposition of such a 3-manifold along tori and Klein bottles, which is the higher dimensional antecedent of the JSJ decompositions considered in this paper.

When the geometric decomposition has only Seifert fibered pieces the vertices are decorated by the quasi-isometry type of the universal cover of the corresponding Seifert fibered manifold. There are only two possible quasi-isometry types, according to whether or not the Seifert fibered piece has boundary. Every vertex in the Bass-Serre tree has infinite valence, so each entry of the structure invariant is either 0 or ∞ .

Behrstock and Neumann [2] state their result in terms of 'bi-similarity'⁴ classes of bi-colored graphs. They show that each bi-similarity class is represented by a unique minimal graph, and that two such 3-manifolds are quasi-isometric if and only if the bi-colored Bass-Serre tree of their geometric decompositions have the same representative minimal graph. Their minimal bi-colored graphs carry exactly the same information as the structure invariant of the decorated Bass-Serre tree. One can construct their graph by taking the vertex set to be the stable decoration set \mathcal{O}_s and connecting vertex $\mathcal{O}_s[j]$ to vertex $\mathcal{O}_s[k]$ by an edge if and only if the j, k -entry of \mathcal{S} is ∞ . The vertices of the graph are 'bi-colored' by the projection $\pi_0: \mathcal{O}_s \rightarrow \mathcal{O}$. Conversely, \mathcal{S} can be recovered by

⁴Their meaning of 'similarity' is different than in this paper.

replacing each edge in the graph by infinitely many edges, lifting the bi-coloring to the universal covering tree, and calculating the structure invariant.

The second paper [3] extends their results to cases where the decomposition involves some hyperbolic pieces. The decorations there are more complex.

4.3 Structure invariants for the JSJ tree of cylinders

Combining Proposition 4.7 with Theorem 2.9 and Corollary 2.10 proves:

Theorem 4.8. *If G is a finitely presented one-ended group not commensurable to a surface group, then the structure invariant for the JSJ tree of cylinders is a quasi-isometry invariant of G , with respect to any of the following initial decorations:*

1. Vertex type: rigid, hanging, or cylinder.
2. Vertex type and, if v is rigid, $\llbracket G_v \rrbracket$.
3. Vertex type and, if v is rigid, $\llbracket (G_v, \mathcal{P}_v) \rrbracket$.

Theorem 4.9. *If G is hyperbolic and the JSJ decomposition of G has no rigid vertices then the invariant of Theorem 4.8 is a complete quasi-isometry invariant.*

Later we will prove a more general result, Theorem 8.5, that includes Theorem 4.9 as a special case. A brief sketch of a direct proof of Theorem 4.9 goes like this: Given two hyperbolic groups as in Theorem 4.9, the structure invariants of Theorem 4.8 are equivalent if and only if the groups have isomorphic decorated JSJ trees of cylinders. Since all non-cylindrical vertices are hanging, it follows, using techniques of Behrstock and Neumann [2], that the groups are quasi-isometric if and only if they have isomorphic decorated JSJ trees of cylinders. Details of the last claim can be found in Dani and Thomas [15, Section 4]. The torsion-free case was previously written up in the thesis of Malone [28].

4.4 Examples

4.4.1 An example with symmetry

Consider the graph of groups Γ in Figure 5, for which $T(\Gamma) = \text{Cyl}(G(\Gamma))$. Let $G := G(\Gamma)$ and $T := T(\Gamma)$.




Figure 5: Symmetric example

Take the initial decoration be vertex type, so $\delta_0(r), \delta_0(r') := \text{'rigid'}$; $\delta_0(c), \delta_0(c') := \text{'cylindrical'}$; and $\delta_0(h) := \text{'hanging'}$.

Let us compute the first neighbor refinement: The image of an edge injection into a rigid vertex group is an infinite index subgroup, so rigid vertices of T are infinite

valence, and adjacent only to cylindrical vertices. Since δ_0 does not distinguish between cylindrical vertices, we have:

$$f_{\tilde{r},0}, f_{\tilde{r}',0} := \begin{cases} \text{'cylindrical'} & \mapsto \infty \\ \text{'rigid'} & \mapsto 0 \\ \text{'hanging'} & \mapsto 0 \end{cases}$$

In this example each edge group surjects onto its adjacent cylindrical vertex group, so cylindrical vertices of T have valence 2. One neighbor is rigid and one is hanging, and δ_0 does not distinguish between the rigid vertices, so we have:

$$f_{\tilde{c},0}, f_{\tilde{c}',0} := \begin{cases} \text{'cylindrical'} & \mapsto 0 \\ \text{'rigid'} & \mapsto 1 \\ \text{'hanging'} & \mapsto 1 \end{cases}$$

Finally, the image of an edge injection into a hanging vertex group is an infinite index subgroup, so hanging vertices of T are infinite valence, and adjacent only to cylindrical vertices. Since δ_0 does not distinguish between cylindrical vertices, we have:

$$f_{\tilde{h},0} := \begin{cases} \text{'cylindrical'} & \mapsto \infty \\ \text{'rigid'} & \mapsto 0 \\ \text{'hanging'} & \mapsto 0 \end{cases}$$

The first refinement is therefore:

$$\begin{aligned} \delta_1(\tilde{r}), \delta_1(\tilde{r}') &:= (\text{'rigid'}, f_{\tilde{r},0}) \\ \delta_1(\tilde{c}), \delta_1(\tilde{c}') &:= (\text{'cylindrical'}, f_{\tilde{c},0}) \\ \delta_1(\tilde{h}) &:= (\text{'hanging'}, f_{\tilde{h},0}) \end{aligned}$$

This is a trivial refinement, so δ_0 is stable, and the structure invariant is given in Table 1.

| | 'cylindrical' | 'rigid' | 'hanging' |
|---------------|---------------|---------|-----------|
| 'cylindrical' | 0 | 1 | 1 |
| 'rigid' | ∞ | 0 | 0 |
| 'hanging' | ∞ | 0 | 0 |

Table 1: Structure invariant for symmetric example

Notice that we get the same structure invariant for the group $G' := G(\Gamma')$ defined by the graph of groups in Figure 6. This is expected, as the groups are quasi-isometric. In fact, G is isomorphic to an index 2 subgroup of G' .




Figure 6: Γ'

4.4.2 The example of Section 3.1

Let Γ be the graph of groups of Figure 2, for which $T(\Gamma) = \text{Cyl}(G(\Gamma))$. Let $G := G(\Gamma)$ and $T := T(\Gamma)$. In this example we take the initial decoration to be by vertex type and quasi-isometry type of vertex stabilizer. The vertices of T carry four possible ornaments, $(r, \llbracket \pi_1(M) \rrbracket)$, $(r, \llbracket F_2 \rrbracket)$, $(h, \llbracket F_2 \rrbracket)$, and $(c, \llbracket \mathbb{Z} \rrbracket)$, where in the first coordinate r , h , and c stand for ‘rigid’, ‘hanging’, and ‘cylindrical’, respectively.

In the first refinement step, observe that each rigid and hanging vertex of T is adjacent to infinitely many cylindrical vertices. The initial decoration does not distinguish any of the cylindrical vertices, so δ_1 does not give any new information about rigid and hanging vertices.

On the other hand, we do distinguish some of the cylindrical vertices. This is because δ_0 does distinguish between three different kinds of rigid and hanging vertex, so we can distinguish cylindrical vertices according to the number and kind of their neighbors.

| | | $(r, \llbracket \pi_1(M) \rrbracket)$ | $(r, \llbracket F_2 \rrbracket)$ | $(h, \llbracket F_2 \rrbracket)$ | $(c, \llbracket \mathbb{Z} \rrbracket)$ |
|--|---|---------------------------------------|----------------------------------|----------------------------------|---|
| \tilde{r}_6 | $(r, \llbracket \pi_1(M) \rrbracket)$ | 0 | 0 | 0 | ∞ |
| $\tilde{r}_1, \tilde{r}_2, \tilde{r}_3, \tilde{r}_4, \tilde{r}_5, \tilde{r}_7, \tilde{r}_8, \tilde{r}_9$ | $(r, \llbracket F_2 \rrbracket)$ | 0 | 0 | 0 | ∞ |
| \tilde{h} | $(h, \llbracket F_2 \rrbracket)$ | 0 | 0 | 0 | ∞ |
| \tilde{c}_1 | $(c, \llbracket \mathbb{Z} \rrbracket)$ | 0 | 5 | 0 | 0 |
| \tilde{c}_2 | | 1 | 1 | 0 | 0 |
| $\tilde{c}_3, \tilde{c}_4, \tilde{c}_6$ | | 0 | 2 | 0 | 0 |
| \tilde{c}_5 | | 0 | 1 | 1 | 0 |

Table 2: First refinement

Table 2 shows δ_1 . In the first column, for convenience, we give representatives of G -orbits of vertex v in T with a given ornament in \mathcal{O}_1 . The second column contains $\delta_0(v)$. The remaining columns in each row give a row vector encoding $f_{v,0}$. For instance, the last row says that a vertex $v \in G\tilde{c}_5$ has $\delta_0(v) = (c, \llbracket \mathbb{Z} \rrbracket)$ and is adjacent to one vertex in $\delta_0^{-1}((r, \llbracket F_2 \rrbracket))$ and one vertex in $\delta_0^{-1}((h, \llbracket F_2 \rrbracket))$.

Table 3 shows δ_2 . The $(c, \llbracket \mathbb{Z} \rrbracket)$ columns are listed in the same order as the $(c, \llbracket \mathbb{Z} \rrbracket)$ rows of Table 2. For instance, the fourth row says that for any given vertex in $G\tilde{r}_2$, $G\tilde{r}_3$, or $G\tilde{r}_5$ in T , it has infinitely many neighbors in $G\tilde{c}_1$ and infinitely many neighbors in $G\tilde{c}_3 \cup G\tilde{c}_4 \cup G\tilde{c}_6$ and no other neighbors.

We can again count neighbors according to the new decoration, but we see in Table 4 that this does not distinguish any additional vertices. Thus, the refinement process has stabilized, and Table 4 is the structure invariant. Up to permuting the ordering on \mathcal{O}_0 and the ordering of rows and columns within each \mathcal{O}_0 -block, this structure invariant

| | | $(r, \llbracket \pi_1(M) \rrbracket)$ | $(r, \llbracket F_2 \rrbracket)$ | $(h, \llbracket F_2 \rrbracket)$ | $(c, \llbracket \mathbb{Z} \rrbracket)$ | | | |
|---|---|---------------------------------------|----------------------------------|----------------------------------|---|----------|----------|----------|
| \tilde{r}_6 | $(r, \llbracket \pi_1(M) \rrbracket)$ | 0 | 0 | 0 | 0 | ∞ | 0 | 0 |
| \tilde{r}_1 | | 0 | 0 | 0 | ∞ | ∞ | 0 | 0 |
| \tilde{r}_4 | $(r, \llbracket F_2 \rrbracket)$ | 0 | 0 | 0 | ∞ | 0 | 0 | ∞ |
| $\tilde{r}_2, \tilde{r}_3, \tilde{r}_5$ | | 0 | 0 | 0 | ∞ | 0 | ∞ | 0 |
| $\tilde{r}_7, \tilde{r}_8, \tilde{r}_9$ | | 0 | 0 | 0 | 0 | 0 | ∞ | 0 |
| \tilde{h} | $(h, \llbracket F_2 \rrbracket)$ | 0 | 0 | 0 | 0 | 0 | 0 | ∞ |
| \tilde{c}_1 | | 0 | 5 | 0 | 0 | 0 | 0 | 0 |
| \tilde{c}_2 | $(c, \llbracket \mathbb{Z} \rrbracket)$ | 1 | 1 | 0 | 0 | 0 | 0 | 0 |
| $\tilde{c}_3, \tilde{c}_4, \tilde{c}_6$ | | 0 | 2 | 0 | 0 | 0 | 0 | 0 |
| \tilde{c}_5 | | 0 | 1 | 1 | 0 | 0 | 0 | 0 |

Table 3: Second refinement

| | | $(r, \llbracket \pi_1(M) \rrbracket)$ | $(r, \llbracket F_2 \rrbracket)$ | $(h, \llbracket F_2 \rrbracket)$ | $(c, \llbracket \mathbb{Z} \rrbracket)$ | | | |
|---|---|---------------------------------------|----------------------------------|----------------------------------|---|----------|----------|----------|
| \tilde{r}_6 | $(r, \llbracket \pi_1(M) \rrbracket)$ | 0 | 0 0 0 0 | 0 | 0 | ∞ | 0 | 0 |
| \tilde{r}_1 | | 0 | 0 0 0 0 | 0 | ∞ | ∞ | 0 | 0 |
| \tilde{r}_4 | $(r, \llbracket F_2 \rrbracket)$ | 0 | 0 0 0 0 | 0 | ∞ | 0 | 0 | ∞ |
| $\tilde{r}_2, \tilde{r}_3, \tilde{r}_5$ | | 0 | 0 0 0 0 | 0 | ∞ | 0 | ∞ | 0 |
| $\tilde{r}_7, \tilde{r}_8, \tilde{r}_9$ | | 0 | 0 0 0 0 | 0 | 0 | 0 | ∞ | 0 |
| \tilde{h} | $(h, \llbracket F_2 \rrbracket)$ | 0 | 0 0 0 0 | 0 | 0 | 0 | 0 | ∞ |
| \tilde{c}_1 | | 0 | 1 1 3 0 | 0 | 0 | 0 | 0 | 0 |
| \tilde{c}_2 | $(c, \llbracket \mathbb{Z} \rrbracket)$ | 1 | 1 0 0 0 | 0 | 0 | 0 | 0 | 0 |
| $\tilde{c}_3, \tilde{c}_4, \tilde{c}_6$ | | 0 | 0 0 1 1 | 0 | 0 | 0 | 0 | 0 |
| \tilde{c}_5 | | 0 | 0 1 0 0 | 1 | 0 | 0 | 0 | 0 |

Table 4: Structure invariant

completely determines T up to δ_0 -preserving tree isomorphism.

5 A new decoration: Stretch factors

5.1 Relative quasi-isometric rigidity

In this section, we associate to the JSJ tree of a cylinders of a group new quasi-isometry invariants that take into account the metric information carried by the various two-ended edge groups. Indeed, consider an infinite order element of an edge stabilizer. Its image in each of the adjacent vertex groups has some translation length, and the ratio of these translation lengths gives a stretch factor that describes how the amalgamation distorts distance as measured in the vertex groups. This stretch factor clearly depends on the

choice of metrics of the vertex groups, so it is not an intrinsic invariant of the group, and, in general, it is not preserved by quasi-isometries. However, we show that when the vertex groups satisfy an appropriate notion of quasi-isometric rigidity—a notion that is satisfied by many interesting classes of groups—then such stretch factors are indeed quasi-isometry invariants.

Definition 5.1. A finitely generated group G is *quasi-isometrically rigid relative to the peripheral structure \mathcal{P}* , or (G, \mathcal{P}) is quasi-isometrically rigid, if there exists a proper geodesic metric space X with peripheral structure \mathcal{P}' and a quasi-isometry $\mu: (G, \mathcal{P}) \rightarrow (X, \mathcal{P}')$ such that

1. $\mu_*(\mathcal{QI}(G, \mathcal{P}))$ is a uniform subgroup of $\mathcal{CI}((X, \mu(\mathcal{P})))$.
2. If $g \in G$ is an infinite order element fixing an element of \mathcal{P} then $i \mapsto \mu(g^i)$ is a coarse similitude.

The pair (X, \mathcal{P}') is called a *rigid model* for (G, \mathcal{P}) .

Remark. In Proposition 5.7 we prove that 1 implies 2 if G is hyperbolic.

Lemma 5.2. *If (X, \mathcal{P}') is a rigid model for (G, \mathcal{P}) then $\mu' \circ \bar{\mu} \in \text{CIsm}((X, \mathcal{P}'))$ for any $\mu, \mu' \in \text{QIsom}((G, \mathcal{P}), (X, \mathcal{P}'))$.*

This fact motivates the terminology ‘rigid’.

Proof. For $\mu, \mu' \in \text{QIsom}((G, \mathcal{P}), (X, \mathcal{P}'))$ we have $\mu' \circ \bar{\mu} \in \text{QIsom}((X, \mathcal{P}'), (X, \mathcal{P}'))$. For any $\mu \in \text{QIsom}((G, \mathcal{P}), (X, \mathcal{P}'))$ we have $\mathcal{CI}((X, \mathcal{P}')) < \mathcal{QI}((X, \mathcal{P}')) = \mu_*(\mathcal{QI}((G, \mathcal{P})))$, so if $\mu_*(\mathcal{QI}((G, \mathcal{P}))) < \mathcal{CI}((X, \mathcal{P}'))$ then $\mathcal{QI}((X, \mathcal{P}')) = \mathcal{CI}((X, \mathcal{P}'))$. \square

Definition 5.3. Let (X, \mathcal{P}') be a rigid model for (G, \mathcal{P}) , and let g be an infinite order element of G that fixes an element of \mathcal{P} . Define the X -length of g , $\ell_X(g)$, to be the multiplicative constant of the coarse similitude from \mathbb{Z} to X defined by $i \mapsto \mu(g^i)$, where $\mu \in \text{QIsom}((G, \mathcal{P}), (X, \mathcal{P}'))$.

If a positive power g^k of an infinite order element g fixes an element of \mathcal{P} define $\ell_X(g) := \frac{1}{k} \ell_X(g^k)$.

Lemma 5.2 implies that $\ell_X(g)$ is independent of the choice of quasi-isometry $\mu \in \text{QIsom}((G, \mathcal{P}), (X, \mathcal{P}'))$.

We remark in the second case that if $\pi: \mathbb{Z} \rightarrow k\mathbb{Z}$ is a closest point projection then $i \mapsto \mu(g^{\pi(i)})$ is a coarse similitude from \mathbb{Z} to X with multiplicative constant $\frac{1}{k} \ell_X(g^k)$, so this is a sensible definition for $\ell_X(g)$.

5.2 Examples of relative quasi-isometric rigidity

Let G be a finitely presented group. Let \mathcal{H} be a finite collection of two-ended subgroups of G . Let \mathcal{P} be the peripheral structure consisting of distinct coarse equivalence classes of conjugates of elements of \mathcal{H} .

1. If G is quasi-isometric to a space X such that $\mathcal{I}(X) = \mathcal{QI}(X)$ then (G, \mathcal{P}) is rigid. The peripheral structure plays no role in this case. Examples include:
 - (a) Irreducible symmetric spaces other than real or complex hyperbolic space; thick Euclidean buildings; and products of such [26, 35].
 - (b) The ‘topologically rigid’ hyperbolic groups of Kapovich and Kleiner [25].
 - (c) Certain Fuchsian buildings [7, 46].
 - (d) Mapping class groups of non-sporadic hyperbolic surfaces [1].
2. If G is quasi-isometric to a space X such that $\mathcal{CI}(X) = \mathcal{QI}(X)$ then (G, \mathcal{P}) is rigid. Again, the peripheral structure plays no role in this case. Xie gives an example of a certain solvable Lie group with this property [47].
3. If X is a real or complex hyperbolic space of dimension at least 3 and G is quasi-isometric to X then (G, \mathcal{P}) is quasi-isometrically rigid whenever \mathcal{H} is non-empty, by a theorem of Schwartz [40].
4. If X' is the 3-valent tree, G is quasi-isometric to X' (so G is virtually free), and G does not virtually split over 0 or 2-ended subgroups relative to \mathcal{H} , then (G, \mathcal{P}) is quasi-isometrically rigid [12, 14]. In this case the model space X depends on \mathcal{P} , and is not necessarily isometric to X' .
5. If $X = \mathbb{H}^2$, $\phi: G \rightarrow X$ is a quasi-isometry, and G does not virtually split over 2-ended subgroups relative to \mathcal{H} , then (G, \mathcal{P}) is quasi-isometrically rigid, as follows. A result of Kapovich and Kleiner [25] shows that G has finite index in $\mathcal{QI}((G, \mathcal{P}))$. Therefore, $\mathcal{QI}((G, \mathcal{P}))$ is a finitely generated group quasi-isometric to X . This quasi-isometry induces a cobounded quasi-action of $\mathcal{QI}((G, \mathcal{P}))$ on X . Such a quasi-action is quasi-isometrically conjugate to an isometric action on X , by a theorem of Markovic [29].

The first four cases actually satisfy the following stronger version of quasi-isometric rigidity:

Definition 5.4. We say G is *strongly quasi-isometrically rigid relative to \mathcal{P}* , or (G, \mathcal{P}) is strongly quasi-isometrically rigid, if there is a proper geodesic space X such that if (X, \mathcal{P}') and (X, \mathcal{P}'') are rigid models for (G, \mathcal{P}) , then there is a coarse isometry ϕ of X such that $\phi(\mathcal{P}') = \mathcal{P}''$.

By contrast, the last case only satisfies the weaker version of rigidity.

For a non-example, consider $G := \mathbb{Z}^n$ and $X := \mathbb{R}^n$. For any \mathcal{H} the group $\mathcal{QI}((G, \mathcal{P}))$ contains maps conjugate to homotheties of \mathbb{R}^n . This implies that the multiplicative constants in $\mathcal{QI}((G, \mathcal{P}))$ are unbounded, so $\mathcal{QI}((G, \mathcal{P}))$ cannot be conjugate into some coarse isometry group.

It is also easy to find non-examples of relative rigidity via splittings: If G virtually splits over a zero or two-ended group relative to \mathcal{H} then (G, \mathcal{P}) is not quasi-isometrically

rigid. In such an example there are generalized Dehn twist quasi-isometries preserving \mathcal{P} , powers of which again produce unbounded multiplicative constants in $\mathcal{QI}((G, \mathcal{P}))$.

The previous examples and non-examples naturally lead to the following question:

Question 1. *If G is a hyperbolic group that is not quasi-isometric to \mathbb{H}^2 and does not virtually split over a zero or two-ended subgroup relative to \mathcal{H} , is (G, \mathcal{P}) quasi-isometrically rigid? Strongly quasi-isometrically rigid?*

5.3 Relative quasi-isometric rigidity for hyperbolic groups

In Theorem 5.6 we give a characterization of relative quasi-isometric rigidity for hyperbolic groups. This combines with Proposition 5.7 to show that the first condition of Definition 5.1 implies the second for hyperbolic groups. Theorem 5.6 also provides an alternative viewpoint that may be useful for resolving Question 1.

A space X is called *visual* if there exists an $A \geq 0$ and $x \in X$ such that for every $y \in X$ there exists an A -coarse-geodesic ray starting at x and passing within distance A of y . It follows from [6, Proposition 5.2 and Proposition 5.6] that if X is quasi-isometric to a visual hyperbolic space then X is a visual hyperbolic space.

Definition 5.5. A map $\phi: X \rightarrow Y$ is an (α, M) -power-quasi-symmetric embedding if for all distinct $x, y, z \in X$

$$\frac{d_Y(\phi(x), \phi(z))}{d_Y(\phi(x), \phi(y))} \leq \eta \left(\frac{d_X(x, z)}{d_X(x, y)} \right),$$

where

$$\eta(r) = \begin{cases} Mr^{1/\alpha} & \text{for } 0 < r < 1, \\ Mr^\alpha & \text{for } 1 \leq r. \end{cases}$$

Let $\text{PQS}((X, d_X), (Y, d_Y))$ denote the set of power-quasi-symmetric homeomorphisms, and abbreviate $\text{PQS}(X, d_X) := \text{PQS}((X, d_X), (X, d_X))$, which is a group. If \mathcal{Z} is a collection of subsets of X , define:

$$\text{PQS}(X, d, \mathcal{Z}) := \{\phi \in \text{PQS}((X, d)) \mid \forall Z \in \mathcal{Z}, \phi(Z) \in \mathcal{Z} \text{ and } \exists! Z' \in \mathcal{Z}, \phi(Z') = Z\}$$

Theorem 5.6. *Let G be a non-elementary hyperbolic group with a peripheral structure \mathcal{P} consisting of coarse equivalence classes of conjugates of finitely many two-ended subgroups. Fix, arbitrarily, a word metric d on G , a basepoint $p \in G$, and a visual metric d_∞ on ∂G . The following are equivalent:*

1. (G, \mathcal{P}) is quasi-isometrically rigid.
2. There exists a proper, geodesic, visual hyperbolic space X and a quasi-isometry $\mu: G \rightarrow X$ such that $\mu_*(\mathcal{QI}((G, \mathcal{P})))$ is a uniform subgroup of $\mathcal{CI}((X, \mu(\mathcal{P})))$.
3. There exists a visual metric d'_∞ on ∂G such that $\text{PQS}(\partial G, d_\infty, \partial \mathcal{P})$ is power-quasi-symmetrically conjugate to a uniform subgroup of $\text{BiLip}(\partial G, d'_\infty)$.

Proof. Since G is a visual hyperbolic space, so is X . Thus, equivalence of 1 and 2 is immediate from the definition of relative rigidity.

Item 2 implies item 3 by taking d'_∞ to be a visual metric on $\partial X = \partial G$ and applying [6, Theorem 6.5], which shows that μ extends to a power-quasi-symmetry of ∂G , and a uniform subgroup of $\mathcal{CI}(X)$ extends to a uniform subgroup of $\text{BiLip}(\partial G, d'_\infty)$.

Conversely, there are several ‘hyperbolic cone’ constructions in the literature [6, 8, 11, 21] that take a metric space Z and produce a hyperbolic metric space $\text{Con}(Z)$ such that a visual metric on $\partial \text{Con}(Z)$ recovers Z with the given metric. We take $\text{Con}_r(Z)$ be the ‘truncated hyperbolic approximation with parameter r ’ of Buyalo and Schroeder [11].

Item 3 implies item 2 by taking X to be the hyperbolic cone $\text{Con}_r(\partial G, d'_\infty)$ for r sufficiently small. \square

Let g be an infinite order element of a hyperbolic group G with a fixed word metric. The isometry L_g defined by left multiplication by g has a well defined translation length, which is positive. If $\mu: G \rightarrow X$ is a quasi-isometry such that $\mu_*(L_g)$ is a coarse isometry, it is not true in general that $\mu_*(L_g)$ has a well defined translation length. The next proposition shows that we do still get a positive translation length for $\mu_*(L_g)$ in the special case of relative rigidity.

Proposition 5.7. *Let G be a hyperbolic group and \mathcal{P} a peripheral structure such that there exists a proper geodesic space X and a quasi-isometry $\mu: G \rightarrow X$ such that $\mu_*(\mathcal{QI}((G, \mathcal{P})))$ is a uniform subgroup of $\mathcal{CI}((X, \mu(\mathcal{P})))$. For every infinite order element $g \in G$ the map $i \mapsto \mu(g^i)$ is a coarse similitude.*

Before proving the proposition we need a few lemmas.

Lemma 5.8. *Let X be proper geodesic space quasi-isometric to a non-elementary hyperbolic group. For $i \in \{0, 1\}$, let ϕ_i be a quasi-isometry of X , with $[\phi_0] = [\phi_1] \in \mathcal{QI}(X)$. The distance between ϕ_0 and ϕ_1 is bounded in terms of the quasi-isometry constants of ϕ_0 and ϕ_1 and the constants of X .*

Proof. The following argument is standard, see for instance [44]. We briefly outline the proof.

Since X is quasi-isometric to a visual hyperbolic space, it is a visual hyperbolic space. Furthermore, every point $x \in X$ can be realized as a quasi-center of an ideal geodesic triangle Δ . Since ϕ_0 and ϕ_1 are coarsely equivalent, $\partial\phi_0 = \partial\phi_1$, so $\phi_0(\Delta)$ and $\phi_1(\Delta)$ are ideal quasi-geodesic triangles with the same ideal vertices, and quasi-geodesic constants depending on those of ϕ_0 and ϕ_1 , respectively. The set of quasi-centers of uniformly quasi-geodesic triangles with the same vertices is bounded in terms of the quasi-geodesic constants and the hyperbolicity constant of X , and $\phi_0(x)$ and $\phi_1(x)$ both lie in this set. \square

Corollary 5.9. *If ϕ is an (M, A) -quasi-isometry and ψ is an A' -coarse isometry with $[\phi] = [\psi] \in \mathcal{QI}(X)$ then there is an A'' depending only on M , A , A' and X such that ϕ is an A'' -coarse isometry.*

Lemma 5.10. *Let $\mu: Y \rightarrow X$ be a quasi-isometry between visual hyperbolic spaces. Suppose that ϕ is a loxodromic isometry of Y with translation length τ . Let y_0 be a point on an axis of ϕ , and set $y_i = \phi^i(y_0)$ and $x_i = \mu(y_i)$. Suppose that $\{\mu_*(\phi^i)\}_{i \in \mathbb{Z}}$ are uniform coarse isometries. Then*

$$L := \lim_{i \rightarrow \infty} \frac{d(x_0, x_i)}{i}$$

exists, and there exists an A such that $i \mapsto \mu(\phi^i(y_0))$ is an $(L, 2A)$ -coarse similitude.

Proof. The fact that $\{\mu_*(\phi^i)\}_{i \in \mathbb{Z}}$ are uniform coarse isometries implies that the difference $|d(x_i, x_j) - d(x_0, x_{j-i})|$ is uniformly bounded. Quasi-geodesic stability further implies that $|d(x_0, x_{i+j}) - d(x_0, x_i) - d(x_0, x_j)|$ is uniformly bounded. Let A be the greater of these two bounds.

Suppose $L^+ := \limsup d(x_0, x_i)/i$ and $L^- := \liminf d(x_0, x_i)/i$ are different. Note $L^- > 0$ since $i \mapsto x_i$ is a quasi-geodesic. Take $\epsilon := (L^+ - L^-)/3$. Choose some i such that $d(x_0, x_i)/i < L^- + \epsilon$ and such that $\alpha := \frac{d(x_0, x_i) + 2A}{d(x_0, x_i)} < \sqrt{\frac{2L^+ + L^-}{L^+ + 2L^-}}$. Choose some j such that $d(x_0, x_i)/i > L^+ - \epsilon$ and such that $\frac{q+1}{q} < \sqrt{\frac{2L^+ + L^-}{L^+ + 2L^-}}$, where q is the integer such that $qi \leq j < (q+1)i$.

The previous inequalities, together with the triangle inequality to decompose $d(x_0, x_j)$ along $x_0, x_i, \dots, x_{qi}, x_j$, yields:

$$\begin{aligned} L^+ - \epsilon &< \frac{d(x_0, x_j)}{j} \leq \frac{d(x_0, x_j)}{qi} \leq \frac{(q+1)(d(x_0, x_i) + 2A)}{qi} \\ &= \frac{\alpha(q+1)d(x_0, x_i)}{qi} \leq \frac{\alpha(q+1)}{q} \cdot (L^- + \epsilon) \\ &< \frac{2L^+ + L^-}{L^+ + 2L^-} \cdot \frac{L^+ + 2L^-}{3} = L^+ - \epsilon \end{aligned}$$

This is a contradiction, so $L^+ = L^- = L$.

Suppose there exists an i such that $d(x_0, x_i) < Li - 2A$. Then $L = \lim_{j \rightarrow \infty} \frac{d(x_0, x_{ij})}{ij} < \lim_{j \rightarrow \infty} \frac{Lij}{ij} = L$, which is a contradiction.

If there exists an i such that $d(x_0, x_i) > Li + 2A$, then a similar computation leads to a contradiction. Therefore, $|d(x_0, x_i) - Li| \leq 2A$, which means $i \mapsto x_i = \mu(\phi^i(y_0))$ is an $(L, 2A)$ -coarse similitude. \square

Proof of Proposition 5.7. Let L_g be left multiplication on G by an infinite order element. For all $n \in \mathbb{Z}$ the map L_g^n is an isometry, so $\mu_*(L_g^n)$ is a quasi-isometry whose constants depend only on those of μ . By hypothesis, there exists an A such that for each n there exists an A -coarse isometry $c_n \in \mathcal{CI}(X)$ with $[c_n] = [\mu_*(L_g^n)]$. By Corollary 5.9, there exists an A' independent of n such that $\mu_*(L_g^n)$ is an A' -coarse isometry. Now apply Lemma 5.10. \square

5.4 Stretch factors

Let G be a finitely presented, one-ended group such that $T := \text{Cyl}(G)$ has two-ended edge stabilizers. Let $\Gamma := G \backslash \text{Cyl}(G)$. Let Y be an algebraic tree of spaces for G over T .

Vertices v_0 and v_1 of T that belong to a common cylinder have stabilizer groups that intersect in a virtually cyclic subgroup.

Recall that $\hat{\Delta}$ denotes the modulus of Definition 2.5.

Definition 5.11. Let v_0 and v_1 be distinct quasi-isometrically rigid vertices of T contained in a common cylinder c . For $i \in \{0, 1\}$, choose a rigid model $(X_{v_i}, \mathcal{P}_{v_i})$ for $(G_{v_i}, \mathcal{P}_{v_i}^\Gamma)$. Let $\langle z \rangle$ be an infinite cyclic subgroup of $G_{v_0} \cap G_{v_1}$, and define the *relative stretch* from v_0 to v_1 to be:

$$\text{relStr}(v_0, v_1, (X_{v_0}, \mathcal{P}_{v_0}), (X_{v_1}, \mathcal{P}_{v_1})) := \frac{\ell_{X_{v_1}}(z)}{\ell_{X_{v_0}}(z)}$$

Clearly, $\text{relStr}(v_0, v_1, (X_{v_0}, \mathcal{P}_{v_0}), (X_{v_1}, \mathcal{P}_{v_1}))$ depends on the choices of $(X_{v_i}, \mathcal{P}_{v_i})$. Recall, by Lemma 5.2, it does not depend on the choice of quasi-isometries $(G_{v_0}, \mathcal{P}_{v_0}^\Gamma) \rightarrow (X_{v_0}, \mathcal{P}_{v_0})$ and $(G_{v_1}, \mathcal{P}_{v_1}^\Gamma) \rightarrow (X_{v_1}, \mathcal{P}_{v_1})$.

Lemma 5.12. $\text{relStr}(v_0, v_1, (X_{v_0}, \mathcal{P}_{v_0}), (X_{v_1}, \mathcal{P}_{v_1}))$ does not depend on the choice of $\langle z \rangle < G_{v_0} \cap G_{v_1}$.

Proof. For $i \in \{0, 1\}$, let e_i be an edge on the geodesic in T between v_0 and v_1 , with $\iota(e_i) = v_i$. Let $\langle z_0 \rangle < G_{e_0}$ and $\langle z_1 \rangle < G_{e_1}$ be infinite cyclic subgroups of minimal index. Since G_{e_0} and G_{e_1} are virtually cyclic, $\langle z \rangle$ has finite index in each of them.

$$\frac{\ell_{X_{v_1}}(z)}{\ell_{X_{v_0}}(z)} = \frac{\ell_{X_{v_1}}(z_1) \cdot \frac{[\langle z_1 \rangle : \langle z_1 \rangle \cap \langle z \rangle]}{[\langle z \rangle : \langle z_1 \rangle \cap \langle z \rangle]}}{\ell_{X_{v_0}}(z_0) \cdot \frac{[\langle z_0 \rangle : \langle z_0 \rangle \cap \langle z \rangle]}{[\langle z \rangle : \langle z_0 \rangle \cap \langle z \rangle]}} = \frac{\ell_{X_{v_1}}(z_1)}{\ell_{X_{v_0}}(z_0)} \cdot \frac{[\langle z_1 \rangle : \langle z_0 \rangle \cap \langle z_1 \rangle]}{[\langle z_0 \rangle : \langle z_0 \rangle \cap \langle z_1 \rangle]} = \frac{\ell_{X_{v_1}}(z_1)}{\ell_{X_{v_0}}(z_0)} \cdot \hat{\Delta}(e_0, e_1)$$

The right-hand side is independent of the choice of z_0 and z_1 , since if, say, $\langle z'_0 \rangle$ is another infinite cyclic subgroup of minimal index in G_{e_0} then:

$$\ell_{X_{v_0}}(z'_0) = \ell_{X_{v_0}}(z_0) \cdot \frac{[\langle z_0 \rangle : \langle z_0 \rangle \cap \langle z'_0 \rangle]}{[\langle z'_0 \rangle : \langle z_0 \rangle \cap \langle z'_0 \rangle]} = \ell_{X_{v_0}}(z_0) \quad \square$$

Corollary 5.13. If v_0 , v_1 , and v_2 are quasi-isometrically rigid vertices in a common cylinder then:

$$\begin{aligned} \text{relStr}(v_0, v_1, (X_{v_0}, \mathcal{P}_{v_0}), (X_{v_1}, \mathcal{P}_{v_1})) \cdot \text{relStr}(v_1, v_2, (X_{v_1}, \mathcal{P}_{v_1}), (X_{v_2}, \mathcal{P}_{v_2})) \\ = \text{relStr}(v_0, v_2, (X_{v_0}, \mathcal{P}_{v_0}), (X_{v_2}, \mathcal{P}_{v_2})) \end{aligned}$$

Proposition 5.14. Let ϕ be a quasi-isometry between finitely presented, one-ended groups G and G' whose JSJ trees of cylinders have two-ended edge stabilizers. Let $\Gamma := G \backslash \text{Cyl}(G)$ and $\Gamma' := G' \backslash \text{Cyl}(G')$. Suppose that v_0 and v_1 are distinct quasi-isometrically rigid vertices of $T(\Gamma)$ contained in a cylinder c . Choose rigid models $(X_{v_0}, \mathcal{P}_{v_0})$ and $(X_{v_1}, \mathcal{P}_{v_1})$ for $(G_{v_0}, \mathcal{P}_{v_0}^\Gamma)$ and $(G_{v_1}, \mathcal{P}_{v_1}^\Gamma)$, respectively. Then:

$$\text{relStr}(v_0, v_1, (X_{v_0}, \mathcal{P}_{v_0}), (X_{v_1}, \mathcal{P}_{v_1})) = \text{relStr}(\phi_*(v_0), \phi_*(v_1), (X_{v_0}, \mathcal{P}_{v_0}), (X_{v_1}, \mathcal{P}_{v_1}))$$

Proof. Let Y be an algebraic tree of spaces for Γ , and let Y' be an algebraic tree of spaces for Γ' . For $v \in \{v_0, v_1\}$ choose $\mu_v \in \text{QIsom}((Y_v, \mathcal{P}_v^\Gamma), (X_v, \mathcal{P}_v))$. Note that $\mu_v \circ \bar{\phi}_v \in \text{QIsom}((Y'_{\phi_*(v)}, \mathcal{P}_{\phi_*(v)}^{\Gamma'}), (X_v, \mathcal{P}_v))$. Define:

$$R := \text{relStr}(v_0, v_1, (X_{v_0}, \mathcal{P}_{v_0}), (X_{v_1}, \mathcal{P}_{v_1}))$$

$$R' := \text{relStr}(\phi_*(v_0), \phi_*(v_1), (X_{v_0}, \mathcal{P}_{v_0}), (X_{v_1}, \mathcal{P}_{v_1}))$$

Choose two points x_0 and x_1 in $\mu_{v_0}(Y_{e_0})$ such that $d_{X_{v_0}}(x_0, x_1) \gg 0$. The idea of the proof is to approximate R by a quantity $Q(x_0, x_1)$ depending on x_0 and x_1 , and similarly approximate R' by $Q'(x_0, x_1)$, and then show:

$$R = \lim_{d(x_0, x_1) \rightarrow \infty} Q(x_0, x_1) = \lim_{d(x_0, x_1) \rightarrow \infty} Q'(x_0, x_1) = R'$$

In the following, quantities are ‘coarsely well defined’ if they are well defined up to additive error independent of the choice of x_0 and x_1 .

By construction, Y_{e_0} is a coset of G_{e_0} , and $\langle z_{e_0} \rangle$ is an infinite cyclic subgroup of minimal index in G_{e_0} . Let $g_0 \in G$ such that $Y_{e_0} = g_0 G_{e_0}$. Since $g_0 \langle z_{e_0} \rangle$ is coarsely dense in Y_{e_0} , there exist integers k_i such that $d_{Y_{v_0}}(\mu^{-1}(x_i), g_0 z_{e_0}^{k_i})$ is small.

Y_{e_0} and Y_{e_1} are coarsely equivalent, so closest point projection $\pi: Y_{e_0} \rightarrow Y_{e_1}$ is coarsely well defined. Moreover, since G_{e_0} and G_{e_1} are commensurable, there exist $\epsilon \in \{\pm 1\}$ and $l \in \mathbb{Z}$ such that $\pi(g_0 z_{e_0}^j)$ is bounded distance from $g_1 z_{e_1}^{\epsilon j \hat{\Delta}(e_0, e_1) + l}$, where $z_{e_1}^{\epsilon j \hat{\Delta}(e_0, e_1) + l}$ is to be interpreted as z_{e_1} raised to the greatest integer less than or equal to $\epsilon j \hat{\Delta}(e_0, e_1) + l$. Now we have the following string of relations, where \sim indicates equality up to additive error in the numerator and denominator, independent of x_0 and x_1 .

$$\begin{aligned} R &= \frac{\ell_{X_{v_1}}(g_1 z_{e_1} g_1^{-1})}{\ell_{X_{v_0}}(g_0 z_{e_0} g_0^{-1})} \cdot \hat{\Delta}(e_0, e_1) \\ &\sim \frac{\left(\frac{d_{X_{v_1}}(\mu_{v_1}(g_1 z_{e_1}^{\epsilon k_0 \hat{\Delta}(e_0, e_1) + l}), \mu_{v_1}(g_1 z_{e_1}^{\epsilon k_1 \hat{\Delta}(e_0, e_1) + l}))}{|(\epsilon k_1 \hat{\Delta}(e_0, e_1) + l) - (\epsilon k_0 \hat{\Delta}(e_0, e_1) + l)|} \right)}{\left(\frac{d_{X_{v_0}}(\mu_{v_0}(g_0 z_{e_0}^{k_0}), \mu_{v_0}(g_0 z_{e_0}^{k_1}))}{|k_1 - k_0|} \right)} \cdot \hat{\Delta}(e_0, e_1) \\ &= \frac{d_{X_{v_1}}(\mu_{v_1}(g_1 z_{e_1}^{\epsilon k_0 \hat{\Delta}(e_0, e_1) + l}), \mu_{v_1}(g_1 z_{e_1}^{\epsilon k_1 \hat{\Delta}(e_0, e_1) + l}))}{d_{X_{v_0}}(\mu_{v_0}(g_0 z_{e_0}^{k_0}), \mu_{v_0}(g_0 z_{e_0}^{k_1}))} \\ &\sim \frac{d_{X_{v_1}}(\mu_{v_1}(\pi(g_0 z_{e_0}^{k_0})), \mu_{v_1}(\pi(g_0 z_{e_0}^{k_1})))}{d_{X_{v_0}}(\mu_{v_0}(g_0 z_{e_0}^{k_0}), \mu_{v_0}(g_0 z_{e_0}^{k_1}))} \\ &\sim \frac{d_{X_{v_1}}(\mu_{v_1} \circ \pi \circ \mu_{v_0}^{-1}(x_0), \mu_{v_1} \circ \pi \circ \mu_{v_0}^{-1}(x_1))}{d_{X_{v_0}}(x_0, x_1)} =: Q(x_0, x_1) \end{aligned}$$

We conclude $R = \lim_{d(x_0, x_1) \rightarrow \infty} Q(x_0, x_1)$.

Similarly, if π' is closest point projection from $\phi(Y_{e_0})$ to $\phi(Y_{e_1})$ define:

$$Q'(x_0, x_1) := \frac{d_{X_{v_1}}(\mu_{v_1} \circ \bar{\phi} \circ \pi' \circ \phi \circ \mu_{v_0}^{-1}(x_0), \mu_{v_1} \circ \bar{\phi} \circ \pi' \circ \phi \circ \mu_{v_0}^{-1}(x_1))}{d_{X_{v_0}}(x_0, x_1)}$$

We have $R' = \lim_{d(x_0, x_1) \rightarrow \infty} Q'(x_0, x_1)$. However, since Y_{e_0} and Y_{e_1} are coarsely equivalent, $\bar{\phi} \circ \pi' \circ \phi$ is coarsely equivalent to π , so $Q(x_0, x_1) \sim Q'(x_0, x_1)$. We conclude:

$$R = \lim_{d(x_0, x_1) \rightarrow \infty} Q(x_0, x_1) = \lim_{d(x_0, x_1) \rightarrow \infty} Q'(x_0, x_1) = R' \quad \square$$

5.5 Uniformization

The stretch factors defined in the previous section depend on the choice of rigid model for the vertex groups. We suppress this dependence by choosing models uniformly:

Definition 5.15. Let $\text{QItypes} := \{[(G, \mathcal{P})]\}$ be the set of quasi-isometry classes of finitely presented groups relative to peripheral structures. For each $Q \in \text{QItypes}$ choose a proper, geodesic space Z_Q with peripheral structure \mathcal{P}_Q such that $Q = [(Z_Q, \mathcal{P}_Q)]$. Define $\text{Model}(Q) := (Z_Q, \mathcal{P}_Q)$. If (Z_Q, \mathcal{P}_Q) is quasi-isometrically rigid then we choose (Z_Q, \mathcal{P}_Q) to be a rigid model as in Section 5.1. We choose $Z_{[\mathbb{R}, \mathbb{R}]} = \mathbb{R}$.

Definition 5.16. If v_0 and v_1 are quasi-isometrically rigid vertices in a cylinder, define:

$$\text{relStr}(v_0, v_1) = \text{relStr}(v_0, v_1, \text{Model}(G_{v_0}, \mathcal{P}_{v_0}^\Gamma), \text{Model}(G_{v_1}, \mathcal{P}_{v_1}^\Gamma))$$

5.6 Normalization for unimodular graphs of groups

Suppose that $T := \text{Cyl}(G)$ has two-ended edge stabilizers and c is a unimodular cylinder in T . Suppose that c contains some quasi-isometrically rigid vertices. Unimodularity implies $\{\text{relStr}(v_0, v_1) \mid v_0, v_1 \in c \text{ are qi rigid}\}$ is bounded. Since stretch factors multiply by Corollary 5.13, there exists a quasi-isometrically rigid v_0 such that for every other quasi-isometrically rigid vertex v_1 in c we have $\text{relStr}(v_0, v_1) \geq 1$. Define $\text{relStr}(c, v_1) := \text{relStr}(v_0, v_1)$.

Definition 5.17. Suppose that $T := \text{Cyl}(G)$ has two-ended edge stabilizers. Let e be an edge of T connecting a cylindrical vertex c to a quasi-isometrically rigid vertex v . Define $\text{relStr}(e) := \text{relStr}(c, v)$.

5.7 An example

Recall Example 1.1: Let $G_i := \langle a, b, t \mid \bar{t}u_i t = v_i \rangle$, where u_i and v_i are words in $\langle a, b \rangle$ given below. In each case G_i should be thought of as an HNN extension of $\langle a, b \rangle$ over \mathbb{Z} with stable letter t . Subdividing the edge gives a unimodular JSJ decomposition Γ_i of G_i whose Bass-Serre T_i is equal to its tree of cylinders.

Let $u_0 := a$, $v_0 := ab\bar{a}\bar{b}^2$, $u_1 := ab$, $v_1 := a^2\bar{b}^2$, $u_2 := ab^2$, and $v_2 := a^2\bar{b}$.

The methods of [14] show that for each i , the pair $\{u_i, v_i\}$ is Whitehead minimal, with Whitehead graph equal to the complete graph on 4 vertices, so $\langle a, b \rangle$ is quasi-isometrically rigid relative to the peripheral structure \mathcal{P}_i coming from incident edge groups, and the rigid model space is just the Cayley graph for $\langle a, b \rangle$ with respect to $\{a, b\}$, ie, the 4-valent tree. Furthermore, $\mathcal{QI}(\langle a, b \rangle, \mathcal{P}_i)$ is transitive on \mathcal{P}_i , and $\llbracket \langle a, b \rangle, \mathcal{P}_1 \rrbracket = \llbracket \langle a, b \rangle, \mathcal{P}_2 \rrbracket = \llbracket \langle a, b \rangle, \mathcal{P}_3 \rrbracket$. Therefore, Theorem 4.8 cannot distinguish G_1 , G_2 , and G_3 .

In T_i , V_C is a single orbit, and each vertex $c \in V_C$ has valence two. Fix some $c \in V_C$, and suppose $\text{Rig}(c) = \{e, e'\}$. Let us assume that e is an edge that attaches to $X_{\tau(e)}$ along the image of a conjugate of $\langle v_i \rangle$ and e' is an edge that attaches to $X_{\tau(e')}$ along the image of a conjugate of $\langle u_i \rangle$. The stabilizers of e and e' are equal to the stabilizer of c , which is infinite cyclic. This, together with the fact that the rigid model vertex space is the Cayley graph (tree) for $\langle a, b \rangle$ with respect to $\{a, b\}$, and the fact that each u_i and v_i is cyclically reduced with respect to $\{a, b\}$, means that $\text{Str}(e)$ is the word length of v_i in $\langle a, b \rangle$, and $\text{Str}(e')$ is the word length of u_i in $\langle a, b \rangle$. Thus, $\text{relStr}(e') = 1$ and $\text{relStr}(e) = |v_i|/|u_i| = 5, 2, \text{ or } 1$, as $i = 0, 1$ or 2 , respectively. By Proposition 5.14, no one of these groups is quasi-isometric to the other.

6 Vertex constraints

In this section we assume that G is a one-ended, finitely presented group such that $T := \text{Cyl}(G)$ has two-ended cylinder stabilizers. Let \mathbf{T} be the quotient graph of cylinders, which is therefore a canonical JSJ decomposition of G over two-ended subgroups; recall Section 2.3.4.

We suppose that X is a tree of spaces over T , quasi-isometric to G .

In Section 4 we saw how to decide if two vertices of T are in the same $\text{Aut}(T, \delta)$ -orbit. In this section we would like to restrict further to subgroups of $\text{Aut}(T, \delta)$ induced by $\mathcal{QI}(X)$, or, in the case that G is hyperbolic, by $\text{Homeo}(\partial X)$. We will actually do something that is weaker in the quasi-isometry case, but has the advantage that the same approach works for both quasi-isometries and boundary homeomorphisms. What we do is restrict to elements of $\text{Aut}(T, \delta)$ that at each vertex look like they are induced by a quasi-isometry or boundary homeomorphism of the appropriate vertex space. We also add a compatibility condition below. First we explain the notation.

For $[\phi] \in \mathcal{QI}(X)$ we can choose a representative ϕ that induces an automorphism ϕ_* of T and splits as a tree of quasi-isometries $\phi_v := \phi|_{X_v} \in \mathcal{QI}((X_v, \mathcal{P}_v), (X_{\phi_*(v)}, \mathcal{P}_{\phi_*(v)}))$ over T .

Similarly, if G is hyperbolic, then X is hyperbolic and $\phi \in \text{Homeo}(\partial X)$ induces an automorphism ϕ_* of T and splits as a tree of boundary homeomorphisms $\phi_v := \phi|_{\partial X_v} \in \text{Homeo}((\partial X_v, \partial \mathcal{P}_v), (\partial X_{\phi_*(v)}, \partial \mathcal{P}_{\phi_*(v)}))$ over T .

Since cylinders are two-ended, each edge space is a quasi-line \mathcal{L} in its respective vertex space. Recall this means that there is a controlled embedding Ξ of \mathbb{R} with image \mathcal{L} . In the hyperbolic case Ξ is actually a quasi-isometric embedding, and \mathcal{L} has distinct endpoints at infinity in the boundary of the vertex space containing it. In this case we

define an orientation of \mathcal{L} to be a choice of one of these boundary points, and a boundary homeomorphism of the vertex space that preserves $\partial\mathcal{L}$ is said to be orientation preserving if it fixes $\partial\mathcal{L}$ and orientation reversing if it exchanges the two points of $\partial\mathcal{L}$.

In the quasi-isometry case we know that $\Xi([0, \infty))$ and $\Xi([0, -\infty))$ are not coarsely equivalent. We define an orientation of \mathcal{L} to be a choice of coarse equivalence class of either $\Xi([0, \infty))$ or $\Xi([0, -\infty))$. A quasi-isometry that coarsely preserves \mathcal{L} is said to be orientation preserving on \mathcal{L} if it fixes the coarse equivalence classes of $\Xi([0, \infty))$ and $\Xi([0, -\infty))$, and orientation reversing if it exchanges them.

We seek $\chi \in \text{Aut}(T, \delta)$ such that for every $v \in \mathcal{V}T$ there exists an element $\phi_v \in \text{Map}((X_v, \mathcal{P}_v), (X_{\chi(v)}, \mathcal{P}_{\chi(v)}))$ such that $(\phi_v)_* = \chi|_{\text{lk}(v)}$, subject to the following compatibility condition. In the quasi-isometry case we require that $(\alpha_{\chi(e)} \circ \phi_{\iota(e)}) \circ (\phi_{\tau(e)} \circ \alpha_e)^{-1}$ is orientation preserving on $X_{\overline{\chi(e)}}$. In the boundary homeomorphism case we require that $(\partial\alpha_{\chi(e)} \circ \phi_{\iota(e)}) \circ (\phi_{\tau(e)} \circ \partial\alpha_e)^{-1}$ is the identity on $\partial X_{\overline{\chi(e)}}$ for every edge e . For brevity, we say “ $(\alpha_{\chi(e)} \circ \phi_{\iota(e)}) \circ (\phi_{\tau(e)} \circ \alpha_e)^{-1}$ is orientation preserving on $X_{\overline{\chi(e)}}$ ” in both cases.

In the boundary homeomorphism case we conclude, in Theorem 7.1, that such a collection of ϕ_v patch together to give $\phi \in \text{Homeo}(\partial X)$ with $\phi_* = \chi$.

The analogous statement is not true for quasi-isometries. To patch together quasi-isometries we need $\alpha_{\chi(e)} \circ \phi_{\iota(e)}$ and $\phi_{\tau(e)} \circ \alpha_e$ to be coarsely equivalent *as maps*, but we have only assumed that they have coarsely equivalent image sets with the same orientations. We also need to know that the ϕ_v have uniform quasi-isometry constants. These points will be addressed in subsequent sections.

6.1 Partial Orientations

A *partial orientation* ζ of X assigns to each cylindrical vertex space and to each peripheral set in each non-elementary vertex space either an orientation of that space or the value ‘NULL’.

A cylindrical vertex space or peripheral set is said to be ζ -oriented if its ζ value is not ‘NULL’, and ζ -unoriented otherwise.

A cylindrical vertex is said to be ζ -oriented or ζ -unoriented if its vertex space is.

An edge $e \in T$ is said to be ζ -oriented or ζ -unoriented if the corresponding edge space in its incident non-elementary vertex is.

The *sign* of a map ϕ that takes an oriented space A to an oriented space B is 1 if the map is orientation preserving and -1 if it is orientation reversing. For a partial orientation ζ we define $\text{sign}_\zeta \phi$ as usual when A and B are both ζ -oriented, and we define $\text{sign}_\zeta \phi := 0$ if either of them is ζ -unoriented.

One partial orientation, ζ' , *extends* another, ζ , if they agree on all ζ -oriented sets.

6.2 Cylindrical vertices

Definition 6.1. Let ζ be a partial orientation. Let c be a cylindrical vertex. The *orientation imbalance* at c with respect to a decoration $\delta: T \rightarrow \mathcal{O}$ and a partial orientation ζ is the function $\Omega_c^{\delta, \zeta}: \mathcal{O} \rightarrow \mathbb{Z}^\mathcal{O}/\{-1, 1\}$, with the action by coordinate-wise multiplication, defined as follows. Choose an orientation of X_c and for each $e \in \text{lk}(c)$ let $\text{sign } \alpha_e$ denote

the sign of α_e with respect to the chosen orientation of X_c and the ζ -orientation of $X_{\bar{e}}$, which we take to be 0 if e is ζ -unoriented. Define:

$$\Omega_c^{\delta, \zeta}(o) := \left[\sum_{e \in \text{lk}(c) \cap \delta^{-1}(o)} \text{sign } \alpha_e \right]$$

If $\Omega_c^{\delta, \zeta}$ is non-zero we call c an *unbalanced cylinder*.

Taking an equivalence class of function in the definition eliminates the dependence on the arbitrary choice of orientation of X_c .

Proposition 6.2. *Suppose δ and ζ are $\text{Map}(X)$ -invariant. If there exists $\phi \in \text{Map}(X)$ such that ϕ_* fixes a cylindrical vertex c and reverses the orientation of X_c then $\Omega_c^{\delta, \zeta}$ is identically zero.*

Proof. Suppose $o \in \mathcal{O}$ is such that there exists a ζ -oriented edge $e \in \text{lk}(c) \cap \delta^{-1}(o)$. Let $v := \tau(e)$. By $\text{Map}(X)$ -invariance, $\zeta(X_{\overline{\phi_*(e)}}) = \phi_v(\zeta(X_{\bar{e}})) = \alpha_{\phi_*(e)} \circ \phi_c \circ \alpha_e^{-1}(\zeta(X_{\bar{e}}))$. Since ϕ_c is orientation reversing, α_e and $\alpha_{\phi_*(e)}$ have opposite signs. Therefore, $(\phi)_*|_{\text{lk}(c)}$ gives a bijection between edges in $\text{lk}(c) \cap \delta^{-1}(o)$ whose attaching map have positive sign and edges in $\text{lk}(c) \cap \delta^{-1}(o)$ whose attaching map have negative sign. Since this is true for every $o \in \mathcal{O}$ such that $\text{lk}(c) \cap \delta^{-1}(o)$ is non-empty, $\Omega_c^{\delta, \zeta}$ is identically zero. \square

Corollary 6.3. *Suppose δ and ζ are $\text{Map}(X)$ -invariant. If G_c contains an infinite dihedral group then $\Omega_c^{\delta, \zeta}$ is identically zero.*

Proposition 6.4. *Suppose δ and ζ are $\text{Map}(X)$ -invariant. For every $\phi \in \text{Map}(X)$ we have $\Omega_c^{\delta, \zeta} = \Omega_{\phi_*(c)}^{\delta, \zeta}$.*

Proof. Choose some orientation on X_c and $X_{\phi_*(c)}$.

If no edge in $\text{lk}(c) \cap \delta^{-1}(o)$ is ζ -oriented then $\Omega_c^{\delta, \zeta}(o) = 0$, and, by $\text{Map}(X)$ -invariance, the same are true for $\phi_*(c)$.

Now consider $o \in \mathcal{O}$ such that there exists an edge $e \in \text{lk}(c) \cap \delta^{-1}(o)$ such that e is ζ -oriented. Let $v := \tau(e)$. By $\text{Map}(X)$ -invariance, $\phi_*(e) \in \text{lk}(\phi_*(c)) \cap \delta^{-1}(o)$ with $\zeta(X_{\overline{\phi_*(e)}}) = \phi_v(\zeta(X_{\bar{e}})) = \alpha_{\phi_*(e)} \circ \phi_c \circ \alpha_e^{-1}(\zeta(X_{\bar{e}}))$. If ϕ_c is orientation reversing then α_e and $\alpha_{\phi_*(e)}$ have opposite signs, so that:

$$\sum_{\phi_*(e) \in \text{lk}(\phi_*(c)) \cap \delta^{-1}(o)} \text{sign } \alpha_{\phi_*(e)} = - \left(\sum_{e \in \text{lk}(c) \cap \delta^{-1}(o)} \text{sign } \alpha_e \right)$$

If ϕ_c is orientation preserving then α_e and $\alpha_{\phi_*(e)}$ have the same signs, so that:

$$\sum_{\phi_*(e) \in \text{lk}(\phi_*(c)) \cap \delta^{-1}(o)} \text{sign } \alpha_{\phi_*(e)} = \sum_{e \in \text{lk}(c) \cap \delta^{-1}(o)} \text{sign } \alpha_e \quad \square$$

The previous proposition shows we can use cylinder imbalances to distinguish different cylinders. The following lemma shows this holds up under refinement of the decoration.

Lemma 6.5. *Suppose $\delta': T \rightarrow \mathcal{O}'$ is a refinement of δ and ζ' is an extension of ζ . Suppose that the δ -partition of edges of T is finer than the partition into ζ -oriented edges and ζ -unoriented edges. Let c be a cylindrical vertex. If $\Omega_c^{\delta, \zeta}$ is non-zero then $\Omega_c^{\delta', \zeta'}$ is non-zero.*

Let c' be a cylindrical vertex distinct from c . If for every $o \in \mathcal{O}$ there exist ζ -oriented edges in $\text{lk}(c) \cap \delta^{-1}(o)$ if and only if there exist ζ -oriented edges in $\text{lk}(c') \cap \delta^{-1}(o)$ then $\Omega_c^{\delta, \zeta} \neq \Omega_{c'}^{\delta, \zeta}$ implies $\Omega_c^{\delta', \zeta'} \neq \Omega_{c'}^{\delta', \zeta'}$.

Proof. If c is unbalanced then there exists an $o \in \mathcal{O}$ such that $\Omega_c^{\delta, \zeta}(o) \neq 0$, which implies that there are ζ -oriented edges in $\text{lk}(c) \cap \delta^{-1}(o)$. Since the δ -partition of edges of T is finer than the partition into ζ -oriented edges and ζ -unoriented edges, all edges in $\text{lk}(c) \cap \delta^{-1}(o)$ are ζ -oriented. Since ζ' extends ζ , all edges in $\text{lk}(c) \cap \delta^{-1}(o)$ are ζ' -oriented, and since δ' refines δ , we have, with respect to ζ' and some fixed orientation of X_c , that:

$$\sum_{e \in \text{lk}(c) \cap \delta^{-1}(o)} \text{sign } \alpha_e = \sum_{o' \in \delta' \circ \delta^{-1}(o)} \sum_{e \in \text{lk}(c) \cap \delta^{-1}(o')} \text{sign } \alpha_e$$

The left hand side is non-zero, so one of the terms of the outer sum on the right hand side must be non-zero. Thus, $\Omega_c^{\delta', \zeta'}$ is not identically zero.

For the second statement, suppose, for contraposition, that $\Omega_c^{\delta', \zeta'} = \Omega_{c'}^{\delta', \zeta'}$. Having chosen orientations on X_c and $X_{c'}$, there is an $\epsilon \in \pm 1$ such that for all $o' \in \mathcal{O}'$:

$$\sum_{e \in \text{lk}(c) \cap (\delta')^{-1}(o')} \text{sign } \alpha_e = \epsilon \left(\sum_{e \in \text{lk}(c') \cap (\delta')^{-1}(o')} \text{sign } \alpha_e \right)$$

If $o \in \mathcal{O}$ is such that there are no ζ -oriented edges in either $\text{lk}(c) \cap \delta^{-1}(o)$ or $\text{lk}(c') \cap \delta^{-1}(o)$ then:

$$\sum_{e \in \text{lk}(c) \cap \delta^{-1}(o)} \text{sign } \alpha_e = 0 = \epsilon \left(\sum_{e \in \text{lk}(c') \cap \delta^{-1}(o)} \text{sign } \alpha_e \right)$$

Otherwise, by hypothesis, there are ζ -oriented edges in both $\text{lk}(c) \cap \delta^{-1}(o)$ and $\text{lk}(c') \cap \delta^{-1}(o)$. We conclude that $\Omega_c^{\delta, \zeta} = \Omega_{c'}^{\delta, \zeta}$ from the following computation, in which the first and third equalities are from the facts that the δ -partition of edges of T is finer than the partition into ζ -oriented edges and ζ -unoriented edges and that δ' refines δ , and the second equality is from the hypothesis that $\Omega_c^{\delta', \zeta'} = \Omega_{c'}^{\delta', \zeta'}$.

$$\begin{aligned} \sum_{e \in \text{lk}(c) \cap \delta^{-1}(o)} \text{sign } \alpha_e &= \sum_{o' \in \delta' \circ \delta^{-1}(o)} \sum_{e \in \text{lk}(c) \cap \delta^{-1}(o')} \text{sign } \alpha_e \\ &= \sum_{o' \in \delta' \circ \delta^{-1}(o)} \epsilon \left(\sum_{e \in \text{lk}(c') \cap (\delta')^{-1}(o')} \text{sign } \alpha_e \right) = \epsilon \left(\sum_{e \in \text{lk}(c') \cap \delta^{-1}(o)} \text{sign } \alpha_e \right) \quad \square \end{aligned}$$

Given δ and ζ that are both $\text{Map}(X)$ -invariant we define the process of *cylinder refinement* as follows.

1. By passing to the coarsest common refinement, we may assume that the δ -partition of edges of T is finer than the partition into ζ -oriented edges and ζ -unoriented edges. The refined δ is still $\text{Map}(X)$ -invariant.
2. Consider the ζ -unoriented, unbalanced cylinders. It is possible to choose orientations of their vertex spaces so that if c and c' are two such cylinders with $\Omega_c^{\delta, \zeta} = \Omega_{c'}^{\delta, \zeta}$ then for all $o \in \mathcal{O}$ we have $\sum_{e \in \text{lk}(c) \cap \delta^{-1}(o)} \text{sign } \alpha_e = \sum_{e \in \text{lk}(c') \cap \delta^{-1}(o)} \text{sign } \alpha_e$. Extend ζ to ζ' by taking these orientations of the unbalanced cylindrical vertex spaces.
3. If e is a ζ -unoriented edge such that $c := \iota(e)$ is cylindrical and ζ' -oriented, define $\zeta'(X_{\bar{e}}) = \alpha_e(\zeta'(X_c))$.
4. Define $\mathcal{O}' := \mathcal{O} \times \{-1, 0, 1\}$. Define $\delta'(e) := (\delta(e), \text{sign}_{\zeta'}(e))$ for each edge and $\delta'(v) := (\delta(v), 0)$ for each vertex.

Lemma 6.6. *Suppose that δ and ζ are both $\text{Map}(X)$ -invariant and that the δ -partition of edges of T is finer than the partition into ζ -oriented edges and ζ -unoriented edges. Let δ' and ζ' be constructed via cylinder refinement, as above. Then δ' is a $\text{Map}(X)$ -invariant refinement of δ and ζ' is a $\text{Map}(X)$ -invariant extension of ζ . Moreover, the δ' -partition on edges is finer than the partition into ζ' -oriented and ζ' -unoriented edges.*

Proof. Suppose e is a ζ -unoriented edge with $\iota(e) := c$ cylindrical and unbalanced. Then e is ζ' -oriented and $\text{sign}_{\zeta'} \alpha_e = 1$, so $\delta'(e) = (\delta(e), 1)$.

By invariance of δ and ζ , Proposition 6.4, and our choice of orientation on X_c and $X_{\phi_*(c)}$, if $\phi \in \text{Map}(X)$ then $\phi_*(c)$ is unbalanced, and $\phi_*(e)$ is ζ -unoriented but ζ' oriented with $\text{sign}_{\zeta'} \alpha_{\phi_*(e')} = 1$. Moreover, ϕ_c is orientation preserving, so $\phi(\zeta'(X_{\bar{e}})) = \zeta'(X_{\overline{\phi_*(e)}})$. It also means that $\delta'(\phi_*(e)) = (\delta(\phi_*(e)), 1) = (\delta(e), 1) = \delta'(e)$.

Now suppose e is ζ -oriented and $\iota(e) := c$ is cylindrical and unbalanced. Invariance of ζ' on e is inherited from invariance of ζ . From the proof of Proposition 6.4, since ϕ_c is orientation preserving, $\text{sign}_{\zeta'} \alpha_e = \text{sign}_{\zeta'} \alpha_{\phi_*(e)}$. Along with invariance of δ , this gives us $\delta'(e) = \delta'(\phi_*(e))$.

For vertices and remaining edges, $\delta'(t) = (\delta(t), 0) = (\delta(\phi_*(t)), 0) = \delta'(\phi_*(t))$.

For the final claim, suppose e is ζ' -oriented and e' is ζ' -unoriented. Since ζ' extends ζ , e' is also ζ -unoriented. If e is ζ -oriented then $\delta(e) \neq \delta(e')$ because the δ -partition of edges of T is finer than the partition into ζ -oriented edges and ζ -unoriented edges. Thus, $\delta'(e) \neq \delta'(e')$, since δ' refines δ .

If e is ζ -unoriented then $\delta'(e) = (\delta(e), \text{sign}_{\zeta'} \alpha_e)$ and $\delta'(e') = (\delta(e'), 0)$ differ in the second coordinate. \square

Lemma 6.7. *Suppose that δ and ζ are $\text{Map}(X)$ -invariant and stable under cylinder refinement. If c is an unbalanced cylindrical vertex and $o \in \mathcal{O}$ such that $\delta^{-1}(o) \cap \text{lk}(c) \neq \emptyset$ then either every edge in $\delta^{-1}(o) \cap \text{lk}(c)$ has orientation preserving attaching map or every edge in $\delta^{-1}(o) \cap \text{lk}(c)$ has orientation reversing attaching map.*

Proof. Cylinder refinement orients every edge in an unbalanced cylinder and distinguishes edges with orientation preserving attaching map from those with orientation reversing attaching map. \square

6.3 Non-elementary vertices

Given a tree of spaces whose underlying tree is decorated $\delta: T \rightarrow \mathcal{O}$, we get a decoration on the peripheral structure of each vertex space via mapping to the tree and composing with δ .

Throughout this subsection we assume that $\delta: T \rightarrow \mathcal{O}$ is a $\text{Map}(X)$ -invariant decoration and ζ is a $\text{Map}(X)$ -invariant partial orientation.

Define:

$$\mathcal{O}' = \mathcal{O} \times \coprod_{Q \in \text{MapTypes}} \text{Map}(Z_Q, \mathcal{P}_Q) \backslash (\mathcal{O}^{\mathcal{P}_Q} \times (\cup_{x \in \partial \mathcal{P}_Q} x \amalg \{\text{'NULL'}\})^{\mathcal{P}_Q} \times (\mathcal{P}_Q \amalg \{\text{'NULL'}\}))$$

The left action of $\text{Map}(Z_Q, \mathcal{P}_Q)$ is given by $\phi.(\chi, \zeta, e) := (\chi \circ \phi^{-1}, \zeta \circ \phi^{-1}, \phi_*(e))$.

If $e \in T$ is an edge with non-elementary terminus $v := \tau(e)$, $Q = \llbracket (X_v, \mathcal{P}_v) \rrbracket$, $\mu_v \in \text{Map}((X_v, \mathcal{P}_v), (Z_Q, \mathcal{P}_Q))$, and ζ_v is a partial orientation on \mathcal{P}_v , define:

$$\delta'(v) := (\delta(v), \text{Map}(Z_Q, \mathcal{P}_Q).(\delta|_{\mathcal{P}_v} \circ \mu_v^{-1}, \zeta_v \circ \mu_v^{-1}, \text{'NULL'}))$$

$$\delta'(e) := (\delta(e), \text{Map}(Z_Q, \mathcal{P}_Q).(\delta|_{\mathcal{P}_v} \circ \mu_v^{-1}, \zeta_v \circ \mu_v^{-1}, (\mu_v)_*(e)))$$

Note that the image is independent of the choice of $\mu_v \in \text{Map}((X_v, \mathcal{P}_v), (Z_Q, \mathcal{P}_Q))$.

Composition of δ' with projection to the first factor of \mathcal{O}' recovers δ , so δ' is a refinement of δ .

Proposition 6.8. *The refinement δ' of δ defined above is $\text{Map}(X)$ -invariant.*

Proof. Take $\phi \in \text{Map}(X)$. If e is an edge with $v := \tau(e)$ non-elementary, $Q := \llbracket (X_v, \mathcal{P}_v) \rrbracket$, and $\chi := \mu_{\phi_*(v)} \circ \phi_v \circ \mu_v^{-1} \in \text{Map}(Z_Q, \mathcal{P}_Q)$, then:

$$\begin{aligned} \delta'(\phi_*(e)) &= \left(\delta(\phi_*(e)), \text{Map}(Z_Q, \mathcal{P}_Q).(\delta|_{\mathcal{P}_{\phi_*(v)}} \circ \mu_{\phi_*(v)}^{-1}, \zeta_{\phi_*(v)} \circ \mu_{\phi_*(v)}^{-1}, (\mu_{\phi_*(v)})_*(\phi_*(e))) \right) \\ &= \left(\delta(e), \text{Map}(Z_Q, \mathcal{P}_Q).(\delta|_{\mathcal{P}_v} \circ \phi_v^{-1} \circ \mu_{\phi_*(v)}^{-1}, \zeta_v \circ \phi_v^{-1} \circ \mu_{\phi_*(v)}^{-1}, (\mu_{\phi_*(v)})_*(\phi_*(e))) \right) \\ &= \left(\delta(e), \text{Map}(Z_Q, \mathcal{P}_Q).(\delta|_{\mathcal{P}_v} \circ \mu_v^{-1} \circ \chi^{-1}, \zeta_v \circ \mu_v^{-1} \circ \chi^{-1}, (\chi \circ \mu_v)_*(e)) \right) \\ &= \left(\delta(e), \text{Map}(Z_Q, \mathcal{P}_Q).(\delta|_{\mathcal{P}_v} \circ \mu_v^{-1}, \zeta_v \circ \mu_v^{-1}, (\mu_v)_*(e)) \right) \\ &= \delta'(e) \end{aligned}$$

Thus, δ' is $\text{Map}(X)$ -invariant. \square

Proposition 6.9. *For vertices $v, w \in T$, $\delta'(v) = \delta'(w)$ if and only if there exists $\phi \in \text{Map}((X_v, \mathcal{P}_v, \delta, \zeta), (X_w, \mathcal{P}_w, \delta, \zeta))$.*

For edges $e, f \in T$ with $v := \tau(e)$ and $w := \tau(f)$ both non-elementary, $\delta'(e) = \delta'(f)$ if and only if there exists $\phi \in \text{Map}((X_v, \mathcal{P}_v, \delta, \zeta), (X_w, \mathcal{P}_w, \delta, \zeta))$ with $\phi_(e) = f$.*

Proof. We give the proof for edges. The proof for vertices is similar.

Let $Q = \llbracket (X_v, \mathcal{P}_v) \rrbracket$. By definition, $\delta'(e) = \delta'(f)$ if and only if $\delta(e) = \delta(f)$ and there exists $\chi \in \text{Map}(Z_Q, \mathcal{P}_Q)$ such that

1. $\delta|_{\mathcal{P}_v} \circ \mu_v^{-1} \circ \chi^{-1} = \delta|_{\mathcal{P}_w} \circ \mu_w^{-1}$
2. $\zeta_v \circ \mu_v^{-1} \circ \chi^{-1} = \zeta_w \circ \mu_w^{-1}$
3. $(\chi \circ \mu_v)_*(e) = (\mu_w)_*(f)$

Define $\phi := \mu_w^{-1} \circ \chi \circ \mu_v \in \text{Map}((X_v, \mathcal{P}_v), (X_w, \mathcal{P}_w))$. Item 1 is equivalent to $\delta|_{\mathcal{P}_v} \circ \phi^{-1} = \delta|_{\mathcal{P}_w}$. Item 2 is equivalent to $\zeta_v \circ \phi^{-1} = \zeta_w$. Item 3 is equivalent to $\phi_*(e) = f$. \square

Corollary 6.10. *There exists a $\text{Map}(X)$ -invariant extension ζ' of ζ such that for any edge e with $v := \tau(e)$ non-elementary, e is ζ' -unoriented if and only if the stabilizer of $X_{\bar{e}}$ in $\text{Map}(X_v, \mathcal{P}_v, \delta, \zeta)$ contains an infinite dihedral group.*

Proof. If e is ζ -unoriented and the stabilizer of $X_{\bar{e}}$ in $\text{Map}(X_v, \mathcal{P}_v, \delta, \zeta)$ does not contain an infinite dihedral group then define an extension ζ' of ζ on $(\delta')^{-1}(\delta'(e))$ as follows. Choose an orientation of $X_{\bar{e}}$. If f is an edge with $\delta'(f) = \delta'(e)$ then, by Proposition 6.9, there exists $\phi \in \text{Map}((X_v, \mathcal{P}_v, \delta, \zeta), (X_w, \mathcal{P}_w, \delta, \zeta))$ with $\phi_*(e) = f$. This means that f is ζ -unoriented, so we extend ζ by defining $\zeta'(X_{\bar{f}}) := \phi(\zeta'(X_{\bar{e}}))$.

The orientation of $X_{\bar{f}}$ is independent of the choice of ϕ because of the stabilizer condition on $X_{\bar{e}}$. \square

Definition 6.11. Given $\text{Map}(X)$ -invariant decoration δ and partial orientation ζ , the process of *vertex refinement* produces the $\text{Map}(X)$ -invariant δ' and partial orientation ζ' defined above.

6.4 Combining the local restrictions

In this section we have our main technical tools. Theorem 6.12 identifies $\text{Aut}(T, \delta)$ orbits. Theorem 6.13 leverages this information to understand decoration preserving isomorphisms between two different trees. Theorem 6.13 provides a blueprint for the main classification theorems in the next two sections.

Theorem 6.12. *Suppose $\delta: T \rightarrow \mathcal{O}$ is a $\text{Map}(X)$ -invariant decoration and ζ is a $\text{Map}(X)$ -invariant partial orientation. Suppose δ and ζ are stable under neighbor, cylinder, and vertex refinement.*

For edges $e, f \in T$ we have $\delta(e) = \delta(f)$ if and only if there exists $\chi \in \text{Aut}(T, \delta)$ such that

- $\chi(e) = f$
- For every $u \in T$ there exists $\phi_u \in \text{Map}((X_u, \mathcal{P}_u, \delta, \zeta), (X_{\chi(u)}, \mathcal{P}_{\chi(u)}, \delta, \zeta))$, such that $\chi|_{\text{lk}(u)} = (\phi_u)_*$.
- For every edge e' the map $(\alpha_{\chi(e')} \circ \phi_{\iota(e')}) \circ (\phi_{\tau(e')} \circ \alpha_{e'})^{-1}$ is orientation preserving on $X_{\overline{\chi(e')}}$.

Proof. If there exists $\chi \in \text{Aut}(T, \delta)$ such that $\chi(e) = f$ then $\delta(e) = \delta(f)$. Conversely, supposing $\delta(e) = \delta(f)$, we construct χ .

Define $\chi(e) := f$.

By Proposition 6.9, there exists $\phi_{\tau(e)} \in \text{Map}((X_{\tau(e)}, \mathcal{P}_{\tau(e)}, \delta, \zeta), (X_{\tau(f)}, \mathcal{P}_{\tau(f)}, \delta, \zeta))$ with $(\phi_{\tau(e)})_*(e) = f$. Define $\chi|_{\text{lk}(\tau(e))} := (\phi_{\tau(e)})_*$.

Now, suppose that we have χ satisfying the desired properties defined on a subtree T' of T such that every leaf is non-elementary and T' contains every edge incident to every non-leaf. Given an edge e_0 with $c := \iota(e_0) \notin T'$ and $\tau(e_0) \in T'$, we show how to extend χ to $\text{lk}(c)$, satisfying the desired properties. Then, by induction, we can extend χ to all of T .

Let $\chi(e_0) := (\phi_{\tau(e_0)})_*(e_0)$. Define $\phi_c := \alpha_{\chi(e_0)}^{-1} \circ \phi_{\tau(e_0)} \circ \alpha_{e_0}$ so that $(\alpha_{\chi(e_0)} \circ \phi_c) \circ (\phi_{\tau(e_0)} \circ \alpha_{e_0})^{-1}$ is orientation preserving on $X_{\overline{\chi(e_0)}}$.

Case 1: c is unbalanced. Extend χ to $\text{lk}(c)$ by choosing a bijection between $\text{lk}(c) \setminus \{e_0\} \cap \delta^{-1}(o)$ and $\text{lk}(\chi(c)) \setminus \{\chi(e_0)\} \cap \delta^{-1}(o)$ for each $o \in \mathcal{O}$. For each o these sets have the same cardinality by neighbor stability. Since c is unbalanced, cylindrical stability implies that c and all edges in $\text{lk}(c)$ are ζ -oriented, and, for each $o \in \mathcal{O}$, all edges in $\delta^{-1}(o)$ have attaching maps with the same sign; recall Lemma 6.7. By $\text{Map}(X)$ -invariance, the same is true for $\chi(c)$, and for each $e_1 \in \text{lk}(c)$ we have $\text{sign}_\zeta \alpha_{e_1} = \text{sign}_\zeta \alpha_{\chi(e_1)}$.

By Proposition 6.9, there exists

$$\phi_{\tau(e_1)} \in \text{Map}((X_{\tau(e_1)}, \mathcal{P}_{\tau(e_1)}, \delta, \zeta), (X_{\tau(\chi(e_1))}, \mathcal{P}_{\tau(\chi(e_1))}, \delta, \zeta))$$

with $(\phi_{\tau(e_1)})_*(e_1) = \chi(e_1)$. Define $\phi|_{X_{\tau(e_1)}} := \phi_{\tau(e_1)}$ and $\chi|_{\text{lk}(\tau(e_1))} := (\phi_{\tau(e_1)})_*$. By construction, $(\alpha_{\chi(e_1)} \circ \phi_c) \circ (\phi_{\tau(e_1)} \circ \alpha_{e_1})^{-1}$ is orientation preserving on $X_{\overline{\chi(e_1)}}$.

In the balanced cases, choose some orientation of X_c and $X_{\chi(c)}$.

Case 2: c is balanced and $o \in \mathcal{O}$ is such that $\delta^{-1}(o) \cap \text{lk}(c) \neq \emptyset$ consists of ζ -oriented edges. By neighbor stability, the total number, n , of edges in $\text{lk}(c) \cap \delta^{-1}(o)$ is equal to the total number of edges in $\text{lk}(\chi(c)) \cap \delta^{-1}(o)$. Since c is balanced, the number of edges in $\text{lk}(c) \cap \delta^{-1}(o)$ with orientation preserving attaching map is equal to the number of edges in $\text{lk}(c) \cap \delta^{-1}(o)$ with orientation reversing attaching map, so there are $n/2$ of each. Cylinder stability implies $\chi(c)$ is also balanced, so there are $n/2$ edges in $\text{lk}(\chi(c)) \cap \delta^{-1}(o)$ with orientation preserving attaching map and $n/2$ with orientation reversing attaching map.

If $\text{sign} \alpha_{e_0} = \text{sign} \alpha_{\chi(e_0)}$ then ϕ_c is orientation preserving. Define χ on $\delta^{-1}(o) \cap \text{lk}(c) \setminus \{e_0\}$ by choosing any bijection with $\delta^{-1}(o) \cap \text{lk}(\chi(c)) \setminus \{\chi(e_0)\}$ that preserves the signs of the attaching maps.

If $\text{sign} \alpha_{e_0} \neq \text{sign} \alpha_{\chi(e_0)}$ then ϕ_c is orientation reversing. Define χ on $\delta^{-1}(o) \cap \text{lk}(c) \setminus \{e_0\}$ by choosing any bijection with $\delta^{-1}(o) \cap \text{lk}(\chi(c)) \setminus \{\chi(e_0)\}$ that exchanges the signs of the attaching maps.

Extend ϕ and χ as in the previous case.

Case 3: c is balanced and $o \in \mathcal{O}$ is such that $\delta^{-1}(o) \cap \text{lk}(c) \neq \emptyset$ consists of ζ -unoriented edges. By neighbor stability, $\text{lk}(c) \setminus \{e_0\} \cap \delta^{-1}(o)$ and $\text{lk}(\chi(c)) \setminus \{e_0\} \cap \delta^{-1}(o)$ have the same cardinality, and we extend χ by an arbitrary bijection between them. Take $e_1 \in \delta^{-1}(o) \cap \text{lk}(c) \setminus \{e_0\}$. By Proposition 6.9, there exists $\phi'_{\tau(e_1)} \in \text{Map}((X_{\tau(e_1)}, \mathcal{P}_{\tau(e_1)}, \delta, \zeta), (X_{\tau(\chi(e_1))}, \mathcal{P}_{\tau(\chi(e_1))}, \delta, \zeta))$ with $(\phi'_{\tau(e_1)})_*(e_1) = \chi(e_1)$. Since e_1 is ζ -unoriented and ζ is stable under vertex refinement, by Corollary 6.10 there exists an element of $\text{Map}(X_{\tau(e_1)}, \mathcal{P}_{\tau(e_1)}, \delta, \zeta)$ reversing $X_{\bar{e}_1}$. Define $\phi_{\tau(e_1)} := \phi'_{\tau(e_1)}$ if $(\alpha_{\chi(e_1)} \circ \phi_c) \circ (\phi'_{\tau(e_1)} \circ \alpha_{e_1})^{-1}$ is orientation preserving on $X_{\overline{\chi(e_1)}}$, and define $\phi_{\tau(e_1)}$ to be $\phi'_{\tau(e_1)}$ precomposed with a $X_{\tau(e_1)}$ -flip otherwise. Extend χ to $\text{lk}(\tau(e_1))$ by $(\phi_{\tau(e_1)})_*$. \square

Let ζ_0 be the trivial partial orientation on X with constant value ‘NULL’. Let $\delta_0: T \rightarrow \mathcal{O}_0$ be any $\text{Map}(X)$ -invariant decoration of T . Perform neighbor, cylinder, and vertex refinement repeatedly until all three stabilize, and let $\delta: T \rightarrow \mathcal{O}$ be the resulting decoration and ζ the resulting partial orientation.

Now suppose X' is a tree of spaces over T' with finite cylinders and such that every $\phi \in \text{Map}(X')$ splits as a tree of maps over T' . Let ζ'_0 be the trivial partial orientation, and let $\delta'_0: T' \rightarrow \mathcal{O}_0$ be a $\text{Map}(X')$ -invariant decoration of T' . (Note that δ_0 and δ'_0 map to the same set of ornaments!) Let ζ' and δ' be the partial orientation extending ζ'_0 and the decoration refining δ'_0 that result from performing neighbor, cylinder, and vertex refinement repeatedly until all three stabilize.

Recall that the process of cylinder refinement involved choosing $\text{Map}(X)$ -invariant orientations. We will need to account for the fact that these choices can be made differently in X and X' . Let $\xi \in \{-1, 1\}^{\mathcal{O}}$. Define $\xi \cdot \zeta$ to be the partial orientation:

$$\xi \cdot \zeta(X_t) = \begin{cases} \text{‘NULL’} & \text{if } \zeta(X_t) = \text{‘NULL’} \\ \zeta(X_t) & \text{if } \xi(\delta(t)) = 1 \\ \text{opposite of } \zeta(X_t) & \text{if } \xi(\delta(t)) = -1 \end{cases}$$

Theorem 6.13. *With the above notation, the following are equivalent:*

1. *There exists $\chi \in \text{Isom}((T, \delta_0), (T', \delta'_0))$ such that:*
 - (a) *For every vertex $v \in T$ there exists $\phi_v \in \text{Map}((X_v, \mathcal{P}_v), (X'_{\chi(v)}, \mathcal{P}'_{\chi(v)}))$, such that $\chi|_{\text{lk}(v)} = (\phi_v)_*$.*
 - (b) *For every edge $e \in T$ we have $(\alpha_{\chi(e)} \circ \phi_{\iota(e)}) \circ (\phi_{\tau(e)} \circ \alpha_e)^{-1}$ is orientation preserving on $X'_{\overline{\chi(e)}}$.*
2. *There exists a bijection $\beta: \text{Im } \delta \rightarrow \text{Im } \delta'$ and $\xi \in \{-1, 1\}^{\mathcal{O}}$ such that:*
 - (a) $\delta_0 \circ \delta^{-1} = \delta'_0 \circ (\delta')^{-1} \circ \beta$
 - (b) *When the rows and columns of $\mathcal{S}(T', \delta', \mathcal{O}')$ are given the β -induced ordering from $\mathcal{S}(T, \delta, \mathcal{O})$, we have $\mathcal{S}(T, \delta, \mathcal{O}) = \mathcal{S}(T', \delta', \mathcal{O}')$.*
 - (c) *For every $o \in \text{Im } \delta$ such that $\delta^{-1}(o)$ consists of non-elementary vertices there exists (equivalently, for every) $v \in \delta^{-1}(o)$ and $v' \in (\delta')^{-1}(\beta(o))$ so that $\text{Map}((X_v, \mathcal{P}_v, \beta \circ \delta, \xi \cdot \zeta), (X'_{v'}, \mathcal{P}'_{v'}, \delta', \zeta'))$ is nonempty.*

- (d) For every $o \in \text{Im } \delta$ such that $\delta^{-1}(o)$ consists of cylindrical vertices there exists (equivalently, for every) $c \in \delta^{-1}(o)$ and $\chi(c) \in (\delta')^{-1}(\beta(o))$ so that $\Omega_c^{\delta, \xi \cdot \zeta} = \Omega_{\chi(c)}^{\delta', \zeta'} \circ \beta$.

Proof. If item 1 is true then $\delta_0 = \delta'_0 \circ \chi$ and $\zeta_0 = \zeta'_0 \circ \chi$. Perform the same sequence of refinements on δ_0 and δ'_0 . Each time the partial orientation on X is extended by choosing some orientation, push that choice forward to X' using the appropriate ϕ_c or ϕ_v . We get the claims of item 2 with β the identity and ξ the constant map sending \mathcal{O} to 1.

We complete the proof by showing that the hypotheses of item 2 allow us to build a isomorphism $\chi \in \text{Isom}((T, \beta \circ \delta), (T', \delta'))$ and a collection of maps ϕ_v satisfying the conditions of item 1. Condition 2a implies $\chi \in \text{Isom}(T, \delta_0), (T', \delta'_0)$. The construction is along the lines of that in the proof of Theorem 6.12: we inductively construct χ and maps $\phi_v \in \text{Map}((X_v, \mathcal{P}_v, \beta \circ \delta, \xi \cdot \zeta), (X'_{\chi(v)}, \mathcal{P}'_{\chi(v)}, \delta', \zeta))$ with $\chi|_{\text{lk}(v)} = (\phi_v)_*$.

To begin, pick a non-elementary vertex $v_0 \in T$ and a vertex $v'_0 \in (\delta')^{-1}(\beta \circ \delta(v_0))$. Define $\chi(v_0) := v'_0$. By 2c and Theorem 6.12, there exists:

$$\phi_{v_0} \in \text{Map}((X_{v_0}, \mathcal{P}_{v_0}, \beta \circ \delta, \xi \cdot \zeta), (X'_{v'_0}, \mathcal{P}'_{v'_0}, \delta', \zeta'))$$

Define $\phi|_{X_{v_0}} := \phi_{v_0}$ and $\chi|_{\text{lk}(v_0)} := (\phi_{v_0})_*$.

Let e_0 be an edge in $\text{lk}(v_0)$ with cylindrical initial vertex $c := \iota(e_0)$. Define $\phi_c := (\alpha_{e'_0})^{-1} \circ \phi_{v_0} \circ \alpha_{e_0}$.

For the induction step we extend χ to $\text{lk}(c)$. When c is balanced the construction is virtually the same as that of Theorem 6.12, so we omit those cases. The remaining case is that c is unbalanced.

Extend χ to $\text{lk}(c)$ by choosing a bijection between $\text{lk}(c) \setminus \{e_0\} \cap \delta^{-1}(o)$ and $\text{lk}(\chi(c)) \setminus \{\chi(e_0)\} \cap \delta'^{-1}(\beta(o))$ for each $o \in \mathcal{O}$. These sets have the same cardinality by condition 2b. Since c is unbalanced, cylindrical stability implies that c and all edges in $\text{lk}(c)$ are ζ -oriented, and, for each $o \in \mathcal{O}$, all edges in $\delta^{-1}(o)$ have attaching maps with the same sign; recall Lemma 6.7. This implies that for each $o \in \mathcal{O}$, $\Omega_c^{\delta, \xi \cdot \zeta}(o) = \pm \Omega_c^{\delta, \zeta}(o)$, so, in particular $\Omega_c^{\delta, \xi \cdot \zeta}$ is not identically zero. Condition 2d then implies $\chi(c)$ is unbalanced, so $\chi(c)$ and all of the edges in $\text{lk}(\chi(c))$ are ζ' -oriented, and edges with the same ornament have attaching maps with the same sign.

We may choose the orientations on X_c and $X'_{\chi(c)}$ to be those given by $\xi \cdot \zeta$ and ζ' , respectively. Together with condition 2d, this implies there exists $\epsilon \in \pm 1$ such that for all $o \in \mathcal{O}$:

$$\sum_{e \in \text{lk}(c) \cap \delta^{-1}(o)} \text{sign}_{\xi \cdot \zeta} \alpha_e = \epsilon \left(\sum_{e' \in \text{lk}(\chi(c)) \cap (\delta')^{-1}(\beta(o))} \text{sign}_{\zeta'} \alpha'_{e'} \right)$$

Since all the edges with a particular ornament have attaching maps of the same sign, this means that for all $e_1 \in \text{lk}(c)$ we have $\text{sign}_{\xi \cdot \zeta} \alpha_{e_1} = \epsilon \text{sign}_{\zeta'} \alpha'_{\chi(e_1)}$. Therefore, the sign of

$$\alpha'_{\chi(e_1)} \circ \phi_c \circ \alpha_{e_1}^{-1} = \alpha'_{\chi(e_1)} \circ (\alpha'_{\chi(e_0)})^{-1} \circ \phi_{v_0} \circ \alpha_{e_0} \circ \alpha_{e_1}^{-1}$$

on $X_{\bar{e}_1}$ with respect to $\xi \cdot \zeta$ and ζ' is $(\epsilon \cdot \text{sign}_{\xi \cdot \zeta} \alpha_{e_0} \cdot \text{sign}_{\xi \cdot \zeta} \alpha_{e_1})^2 = +1$.

By Proposition 6.9 and condition 2c, there exists

$$\phi_{\tau(e_1)} \in \text{Map}((X_{\tau(e_1)}, \mathcal{P}_{\tau(e_1)}, \beta \circ \delta, \xi \cdot \zeta), (X'_{\tau(\chi(e_1))}, \mathcal{P}'_{\tau(\chi(e_1))}, \delta', \zeta'))$$

with $(\phi_{\tau(e_1)})_*(e_1) = \chi(e_1)$. Define $\phi|_{X_{\tau(e_1)}} := \phi_{\tau(e_1)}$ and $\chi|_{\text{lk}(\tau(e_1))} := (\phi_{\tau(e_1)})_*$. We know $(\alpha_{\chi(e_1)} \circ \phi_c) \circ (\phi_{\tau(e_1)} \circ \alpha_{e_1})^{-1}$ is orientation preserving on $X'_{\chi(e_1)}$ because:

$$\alpha'_{\chi(e_1)} \circ \phi_c \circ \alpha_{e_1}^{-1}(\xi \cdot \zeta(X_{\bar{e}_1})) = \zeta'(X'_{\chi(e_1)}) = \phi_{\tau(e_1)}(\xi \cdot \zeta(X_{\bar{e}_1}))$$

We remark that it is not required that $\phi_c(\xi \cdot \zeta(X_c)) = \zeta'(X'_{\chi(c)})$, but this can easily be arranged by redefining $\xi(\delta(c))$ to be $\epsilon \cdot \xi(\delta(c))$. \square

7 Classification of hyperbolic groups up to boundary homeomorphism from their two-ended JSJ splittings

We are now ready to prove our first classification theorem, characterizing the homeomorphism type of the Gromov boundary of a one-ended hyperbolic group from its JSJ tree of cylinders.

Theorem 7.1. *Let G be a one-ended hyperbolic group with non-trivial JSJ decomposition over two-ended subgroups, with $T := \text{Cyl}(G)$. Let X be an algebraic tree of spaces for G over T . Let ζ_0 be the trivial partial orientation on X . Take the initial decoration δ_0 on T to be by vertex type ('cylindrical', 'rigid', or 'hanging') and relative boundary homeomorphism type. Perform neighbor, cylinder, and vertex refinement until all three stabilize to give a decoration $\delta: T \rightarrow \mathcal{O}$ and a partial orientation ζ of X .*

Let G' be another one-ended hyperbolic group with non-trivial JSJ decomposition over two-ended subgroups. Define T' , X' , δ'_0 , ζ'_0 , $\delta': T' \rightarrow \mathcal{O}'$, and ζ' as we did for G . Then ∂G is homeomorphic to $\partial G'$ if and only if there exists a bijection $\beta: \text{Im } \delta \rightarrow \text{Im } \delta'$ and a $\xi \in \{-1, 1\}^{\mathcal{O}}$ such that:

1. $\delta_0 \circ \delta^{-1} = \delta'_0 \circ (\delta')^{-1} \circ \beta$
2. *When the rows and columns of $\mathcal{S}(T', \delta', \mathcal{O}')$ are given the β -induced ordering from $\mathcal{S}(T, \delta, \mathcal{O})$, we have $\mathcal{S}(T, \delta, \mathcal{O}) = \mathcal{S}(T', \delta', \mathcal{O}')$.*
3. *For every $o \in \text{Im } \delta$ such that $\delta^{-1}(o)$ consists of non-elementary vertices there exists (equivalently, for every) $v \in \delta^{-1}(o)$ and $v' \in (\delta')^{-1}(\beta(o))$ so that*

$$\text{Homeo}((\partial X_v, \partial \mathcal{P}_v, \beta \circ \delta, \xi \cdot \zeta), (\partial X'_{v'}, \partial \mathcal{P}'_{v'}, \delta', \zeta'))$$

is nonempty.

4. *For every $o \in \text{Im } \delta$ such that $\delta^{-1}(o)$ consists of cylindrical vertices there exists (equivalently, for every) $c \in \delta^{-1}(o)$ and $c' \in (\delta')^{-1}(\beta(o))$ such that $\Omega_c^{\delta, \xi \cdot \zeta} = \Omega_{c'}^{\delta', \zeta'} \circ \beta$.*

Proof. Since the initial decorations δ_0 and δ'_0 are trivial, the given conditions are equivalent, by Theorem 6.13 for boundary homeomorphism, to the existence of $\chi \in \text{Isom}(T, T')$ such that:

1. For every vertex $v \in T$ there exists $\phi_v \in \text{Homeo}((\partial X_v, \partial \mathcal{P}_v), (\partial X'_{\chi(v)}, \partial \mathcal{P}'_{\chi(v)}))$, such that $\chi|_{\text{lk}(v)} = (\phi_v)_*$.
2. For every edge $e \in T$ we have $\partial \alpha_{\chi(e)} \circ \phi_{\iota(e)} = \phi_{\tau(e)} \circ \partial \alpha_e$.

These conditions say there exists an isomorphism $\chi: T \rightarrow T'$ and a tree of boundary homeomorphisms over χ compatible with X and X' . By Theorem 2.28, this is equivalent to the existence of a boundary homeomorphism between ∂X and $\partial X'$, hence between ∂G and $\partial G'$. \square

Example 7.2. Recall the example of Section 5.7. The trees of cylinders of these groups are isomorphic: they are bipartite with one orbit of valence 2 cylindrical vertex and one orbit of infinite valence relatively rigid vertex. Recall also that the rigid vertices were discussed in Section 3, and the relative quasi-isometry (hence, boundary homeomorphism) types of these three examples are the same.

There are two G_i -orbits of edge in $\text{Cyl}(G_i)$, so the only remaining question about $\text{Homeo}(\partial G_i)$ -orbits is whether there are one or two orbits of edges. For boundary homeomorphism we do not care about stretch factors on the edges, so the edges have trivial initial decoration.

The quasi-isometry group of the vertex stabilizer preserving the peripheral structure coming from incident edge groups is transitive on peripheral sets, so the vertex refinement is a trivial refinement of the initial decoration.

Moreover, the quasi-isometry group of the vertex stabilizer preserving the peripheral structure coming from incident edge groups contains an infinite dihedral group in the stabilizer of each peripheral subset, so all cylinders are balanced, and the cylinder refinement is a trivial refinement of the initial decoration.

We conclude there is only one $\text{Homeo}(\partial G_i)$ orbit of edges in $\text{Cyl}(G_i)$, so the initial decoration and initial (trivial) partial orientation are stable. Since the structure invariants are the same for the three G_i , they have homeomorphic boundaries.

8 Quasi-isometry classification of groups from their two-ended JSJ splittings

We are now almost ready to prove our second main theorem, characterizing the quasi-isometry type of a finitely presented one-ended group from its JSJ tree of cylinders. Before doing so, we explain the extreme flexibility provided by the hanging vertices of the tree.

8.1 Quasi-isometric flexibility of hanging spaces

Recall that the fixed model space for hanging vertices is the universal cover of a fixed hyperbolic pair of pants Σ , with peripheral structure consisting of the coarse equivalence classes of the boundary components of $\tilde{\Sigma}$.

Proposition 8.1 (cf [2, Theorem 1.2]). *Let G be a finitely presented, one-ended group admitting a JSJ decomposition over two-ended subgroups with two-ended cylinder stabilizers. Let $\Gamma := G \backslash \text{Cyl}(G)$ and $T := \text{Cyl}(G) = T(\Gamma)$. Let X be an algebraic tree of spaces for G over T . Let v be a hanging vertex in Γ . Let $\delta: T \rightarrow \mathcal{O}$ be a $\mathcal{QI}(X)$ -invariant decoration and let ζ be a $\mathcal{QI}(X)$ -invariant partial orientation. For each edge $e \in \Gamma$ incident to v , choose a positive real parameter σ_e . Let $\delta': \partial\tilde{\Sigma} \rightarrow \mathcal{O}$ be a decoration of the peripheral structure of $\tilde{\Sigma}$ and let ζ' be a partial orientation of $\partial\Sigma$ such that $\mathcal{QI}(\tilde{\Sigma}, \partial\tilde{\Sigma}, \delta', \zeta')$ acts coboundedly on $\tilde{\Sigma}$.*

Suppose that for some $e_0 \in \text{lk}(\tilde{v})$ we are given a coarse similitude $\phi|_{X_{e_0}}$ from X_{e_0} to a component B_0 of $\partial\tilde{\Sigma}$ that respects the decoration and partial orientations. Suppose further that there exists a $\phi' \in \text{QIsom}((X_v, \mathcal{P}_v, \delta, \zeta), (\tilde{\Sigma}, \partial\tilde{\Sigma}, \delta', \zeta'))$ such that $\phi'(X_e) = B_0$. Then there exists $\phi \in \text{QIsom}((X_v, \mathcal{P}_v, \delta, \zeta), (\tilde{\Sigma}, \partial\tilde{\Sigma}, \delta', \zeta'))$ extending $\phi|_{X_{e_0}}$ that, for each edge $e \in \text{lk}(v) \setminus \{e_0\}$, restricts to be a coarse similitude with multiplicative constant σ_e on X_e .

The quasi-isometry constants of ϕ can be bounded in terms of G_v , \mathcal{P}_v , the coboundness constant for $\mathcal{QI}(\tilde{\Sigma}, \partial\tilde{\Sigma}, \delta', \zeta') \curvearrowright \tilde{\Sigma}$, the constants of $\phi|_{X_{e_0}}$, and the σ_v .

Sketch. The proof follows the same argument as [2, Theorem 1.2]. The idea is to build ϕ inductively, peripheral set by peripheral set. We start with $\phi_{X_{e_0}}$. Let σ_0 be the multiplicative constant of $\phi_{X_{e_0}}$. Then we want to extend ϕ to peripheral sets that come close to X_{e_0} in X_v , sending these to components of $\partial\tilde{\Sigma}$ that come close to $\phi(X_{e_0})$.

For another peripheral set X_e , the number of peripheral sets in the $\mathcal{QI}(X_v, \mathcal{P}_v, \delta, \zeta)$ -orbit of X_e that come within some fixed distance K of a subsegment of X_{e_0} of length l is coarsely dl for some $d > 0$.

Let d'_r be such that there are coarsely $d'_r l$ peripheral sets in the $\mathcal{QI}(\tilde{\Sigma}, \partial\tilde{\Sigma}, \delta', \zeta')$ -orbit of $\phi'(X_e)$ that come within r of a subsegment of B_0 of length l . The fact that $\mathcal{QI}(\tilde{\Sigma}, \partial\tilde{\Sigma}, \delta', \zeta') \curvearrowright \tilde{\Sigma}$ is C -cobounded for some C says that $d'_C > 0$, and, in fact, d'_r grows exponentially in r . This means that there is a logarithmically growing function whose value R at $\frac{d}{\sigma_0}$ is such that $d'_R \geq \frac{d}{\sigma_0}$. Thus, for any l there is a way to send the dl elements in the $\mathcal{QI}(X_v, \mathcal{P}_v, \delta, \zeta)$ -orbit of X_e that come within distance K of a length l subsegment S of X_{e_0} injectively to the peripheral sets in the $\mathcal{QI}(\tilde{\Sigma}, \partial\tilde{\Sigma}, \delta', \zeta')$ -orbit of $\phi'(X_e)$ that come within R of the length approximately $\sigma_0 l$ subsegment $\phi|_{X_e}(S)$ of B_0 .

In this way one builds a matching between the peripheral sets that come close to X_{e_0} and the peripheral sets that come close to B_0 , respecting decorations and partial orientations. Then ϕ is defined along such a matched pair to be a coarse similarity with the appropriate σ_e as multiplicative constant, and the neighbor-matching is repeated for each such pair. \square

8.2 Geometric trees of spaces for groups with two-ended cylinders stabilizers

Let G be a finitely presented, one-ended group admitting a JSJ decomposition over two-ended subgroups with two-ended cylinder stabilizers. Let $\Gamma := G \backslash \text{Cyl}(G)$ and $T := \text{Cyl}(G) = T(\Gamma)$.

Recall that in Section 2.5 we built an algebraic tree of spaces Y quasi-isometric to G , and gave conditions for a collection of quasi-isometries of the vertex spaces to patch together to give a quasi-isometry of Y . Now we will construct a *geometric tree of spaces* X by uniformizing the vertex spaces, that is, replacing each vertex space by its uniform model from Section 5.5. The quasi-isometries between vertex spaces and their uniform models will patch together to give a quasi-isometry from Y to X . Therefore, X will be quasi-isometric to G . The price to pay for uniformizing the vertex spaces is that in general G only admits a cobounded quasi-action on X , not a cocompact action, but this will not affect us.

We use the same notation as in Section 2.5. Let Y be the algebraic tree of spaces constructed there. For a relatively rigid vertex $v \in \Gamma$, fix a quasi-isometry $\nu_v: (G_v, \mathcal{P}_v) \rightarrow (Z_{\llbracket(G_v, \mathcal{P}_v)\rrbracket}, \mathcal{P}_{\llbracket(G_v, \mathcal{P}_v)\rrbracket})$ from G_v to the chosen model space for the relative quasi-isometry type of (G_v, \mathcal{P}_v) .

If G_v is virtually cyclic choose a cyclic subgroup $\langle z_v \rangle < G_v$ of minimal index. Define ν_v by sending G_v onto $\langle z_v \rangle$ by closest point projection and sending z_v^k to $k\sigma_v \in \mathbb{R}$, where σ_v is a positive real parameter chosen as follows. If v is not adjacent to any quasi-isometrically rigid vertices then choose $\sigma_v := 1$. Otherwise, choose $\sigma_v := \min \ell_{X_w}(z_v)$, where the minimum is taken over quasi-isometrically rigid vertices w adjacent to v .

Remark 8.2. This choice of σ_v 's is convenient because it will imply, for an edge e with $\iota(e)$ cylindrical and $\tau(e)$ rigid, that the attaching map α_e^X constructed below is a coarse similitude whose multiplicative constant is equal to the stretch factor $\text{relStr}(e)$ defined in Section 5.6.

For a hanging vertex $v \in \Gamma$ we define ν_v as follows. For each edge $e \in \text{lk}(v)$ choose a cyclic subgroup $\langle z_e \rangle < G_e$ of minimal index. Define

$$\nu_v: (G_v, \mathcal{P}_v) \rightarrow (Z_{\llbracket(G_v, \mathcal{P}_v)\rrbracket}, \mathcal{P}_{\llbracket(G_v, \mathcal{P}_v)\rrbracket})$$

to be a quasi-isometry such that for each $e \in \text{lk}(v)$, each coset of G_e , which is a peripheral set in \mathcal{P}_v , is sent to a peripheral set in $\mathcal{P}_{\llbracket(G_v, \mathcal{P}_v)\rrbracket}$ by a coarse similitude with multiplicative constant:

$$\frac{[\langle z_{\tau(e)} \rangle : \langle z_{\tau(e)} \rangle \cap \langle z_e \rangle]}{[\langle z_e \rangle : \langle z_{\tau(e)} \rangle \cap \langle z_e \rangle]} \cdot \sigma_{\tau(e)} \\ \ell_{G_v}(z_e)$$

Here, $\ell_{G_v}(z_e)$ is the translation length of z_e in the Cayley graph of G_v , which is non-zero since G_v is hyperbolic, and $\sigma_{\tau(e)}$ is the parameter for $G_{\tau(e)}$ chosen above. Such a quasi-isometry can be constructed using Proposition 8.1. These particular values are chosen to make Lemma 8.3, below, true.

Now, for each vertex $v \in \mathcal{VT}$ define X_v to be a copy of $Z_{[(G_v, \mathcal{P}_v)]}$ with isometry $\mu_v: X_v \rightarrow Z_{[(G_v, \mathcal{P}_v)]}$. Define $\phi_v: Y_v \rightarrow X_v$ by $x \mapsto \mu_v^{-1} \circ \nu_v(h_{(v,i)}^{-1}x)$.

We define edge spaces and attaching maps in X to be compatible with those of Y , as follows. Consider an edge e with $v := \iota(e)$ and $w := \tau(e)$. There are $h_{(v,i)}$ and $g_{(e,j)}$ such that $v = h_{(v,i)}\tilde{v}$ and $e = h_{(v,i)}g_{(e,j)}\tilde{e}$. Define $\alpha_e^X := \phi_w \circ \alpha_e^Y \circ \pi_{Y_e} \circ \phi_v^{-1}$, where π_{Y_e} denotes closest point projection to Y_e . The map α_e^X is coarsely well defined, since π_{Y_e} moves points of $\phi_v^{-1}(X_e)$ bounded distance. Define $\alpha_{\bar{e}}^X := \phi_v \circ \alpha_{\bar{e}}^Y \circ \pi_{Y_{\bar{e}}} \circ \phi_w^{-1}$, where $\pi_{Y_{\bar{e}}}$ is closest point projection from Y_w to the coarsely dense subset $Y_{\bar{e}}$. This map is well defined, and is a quasi-isometry inverse to α_e^X , since $\pi_{Y_{\bar{e}}}$ moves points bounded distance.

Chasing through these definitions on easily demonstrates:

Lemma 8.3. *If $c = \iota(e)$ is cylindrical and $v = \tau(e)$ is hanging then $\alpha_e^X: X_c \rightarrow X_v$ is a coarse isometric embedding.*

Proposition 8.4. *With notation as above, G , X , and Y are quasi-isometric to one another.*

Proof. G is quasi-isometric to Y by Lemma 2.15. Proposition 2.16 implies X and Y are quasi-isometric, since (ϕ_v) is a tree of quasi-isometries over the identity on T compatible with X and Y . \square

8.3 Quasi-isometries

Let G be a finitely presented, one-ended group with non-trivial JSJ decomposition over two-ended subgroups such that:

- Every non-elementary vertex is either hanging or quasi-isometrically rigid relative to the peripheral structure coming from incident edge groups.
- Cylinder stabilizers are two-ended.

If Question 1 has positive answer then every one-ended hyperbolic group with a non-trivial JSJ decomposition is of this form.

Theorem 8.5. *Let G and G' be finitely presented, one-ended groups with non-trivial JSJ decompositions over two-ended subgroups with two-ended cylinder stabilizers.*

Let $T := \text{Cyl}(G)$. Let X be a geometric tree of spaces for G over T , as in Section 8.2. Let ζ_0 be the trivial partial orientation on X . Let δ_0 the decoration on T that sends an edge e incident to a rigid vertex to its relative stretch factor $\text{relStr}(e)$ as in Definition 5.17, sends other edges to ‘NULL’, and sends vertices to their vertex type (‘cylindrical’, ‘rigid’, or ‘hanging’) and relative quasi-isometry type. Let ζ_0 be the trivial partial orientation on X . Perform neighbor, cylinder, and vertex refinement until all three stabilize to give a decoration $\delta: T \rightarrow \mathcal{O}$ and a partial orientation ζ of X .

Define T' , X' , δ'_0 , ζ'_0 , $\delta': T' \rightarrow \mathcal{O}'$, and ζ' for G' as we did for G . In particular, X' is uniformized with respect to the same choice of model spaces from Definition 5.15.

Then G and G' are quasi-isometric if and only if there exists a bijection $\beta: \text{Im } \delta \rightarrow \text{Im } \delta'$ and $\xi \in \{-1, 1\}^{\mathcal{O}}$ such that:

1. $\delta_0 \circ \delta^{-1} = \delta'_0 \circ (\delta')^{-1} \circ \beta$
2. When the rows and columns of $\mathcal{S}(T', \delta', \mathcal{O}')$ are given the β -induced ordering from $\mathcal{S}(T, \delta, \mathcal{O})$, we have $\mathcal{S}(T, \delta, \mathcal{O}) = \mathcal{S}(T', \delta', \mathcal{O}')$.
3. For every $o \in \text{Im } \delta$ such that $\delta^{-1}(o)$ consists of non-elementary vertices there exists (equivalently, for every) $v \in \delta^{-1}(o)$ and $v' \in (\delta')^{-1}(\beta(o))$ so that

$$\text{QIsom}((X_v, \mathcal{P}_v, \beta \circ \delta, \xi \cdot \zeta), (X'_{v'}, \mathcal{P}'_{v'}, \delta', \zeta'))$$

is nonempty.

4. For every $o \in \text{Im } \delta$ such that $\delta^{-1}(o)$ consists of cylindrical vertices, there exists (equivalently, for every) $c \in \delta^{-1}(o)$ and $c' \in (\delta')^{-1}(\beta(o))$ such that $\Omega_c^{\delta, \xi \cdot \zeta} = \Omega_{c'}^{\delta', \zeta'} \circ \beta$.

The construction is a modification of the proof of Theorem 6.13. Recall in that case we inductively built $\chi \in \text{Isom}((T, \delta), (T, \delta'))$ and quasi-isometries

$$\phi_v \in \text{QIsom}((X_v, \mathcal{P}_v, \beta \circ \delta, \xi \cdot \zeta), (X'_{\chi(v)}, \mathcal{P}'_{\chi(v)}, \delta', \zeta'))$$

such that $(\phi_v)_* = \chi|_{\text{lk}(v)}$. The proof of Theorem 6.13 mainly focuses on the inductive step in the link of a cylindrical vertex, and chooses any ϕ_v as above such that $(\phi_v)_*$ agrees with χ on the incoming edge to v .

In the present context we must be more careful about the choices of the ϕ_v . The proof in Theorem 6.13 gives us a collection of quasi-isometries (ϕ_v) such that for every edge $e \in T$ with $\iota(e)$ cylindrical we have that $(\alpha_{\chi(e)} \circ \phi_{\iota(e)}) \circ (\phi_{\tau(e)} \circ \alpha_e)^{-1}$ is orientation preserving on $X'_{\chi(e)}$, but now we require it to be coarsely the identity on $X'_{\chi(e)}$. Furthermore, we need the quasi-isometry constants of the ϕ_v to be uniformly bounded.

Here is how we achieve these requirements. Hanging vertices present no obstacles, since by Proposition 8.1 they are so flexible. The real work is in dealing with the rigid vertices. For these we choose in advance a finite number of quasi-isometries to use as building blocks. Since the collection is finite, the constants are uniformly bounded. We will choose the maps on cylinder spaces to be coarse isometries. It then remains to see that if $e \in \mathcal{ET}$ is an edge with $c := \iota(e)$ cylindrical and $v := \tau(e)$ relatively rigid, that we can make ϕ_c agree with a map ϕ_v constructed from the pre-chosen building blocks. We assume we have chosen enough building blocks so that we can make $(\phi_v)_*(e) = (\phi_c)_*(e)$, with the correct orientation on $X_{\bar{e}}$. This is handled by the same considerations as Theorem 6.13. Additionally, we have set up the geometric tree of spaces so that the edge inclusion into a rigid vertex is a coarse similitude whose multiplicative constant is the stretch factor of the edge. Since we have incorporated the stretch factors into the decorations, we are guaranteed that the stretch factor on e matches the stretch factor on $\chi(e)$. It follows that $(\alpha_{\chi(e)} \circ \phi_c) \circ (\phi_v \circ \alpha_e)^{-1}$ is a coarse isometry that is orientation preserving. Finally, we make it coarsely the identity by adjusting ϕ_v using the group action.

Proof of Theorem 8.5. By Theorem 2.9, Corollary 2.10, Proposition 2.16, and Proposition 5.14, X and X' are quasi-isometric if and only if there exists a tree of quasi-isometries over an element of $\text{Isom}((T, \delta_0), (T', \delta'_0))$ compatible with X and X' .

The existence of a tree of quasi-isometries over an element of $\text{Isom}((T, \delta_0), (T', \delta'_0))$ compatible with X and X' implies the above conditions. Our goal is to show the converse.

Suppose $o \in \mathcal{O}$ is an ornament such that $\delta^{-1}(o)$ consists of vertices that are relatively quasi-isometrically rigid. Choose representatives $v_{o,1}, \dots, v_{o,i_o}$ of the G -orbits contained in $\delta^{-1}(o)$. Suppose $o' \in \mathcal{O}$ is an ornament such that $\delta^{-1}(o')$ consists of edges incident to $v \in \delta^{-1}(o)$. For each $1 \leq i \leq i_o$ choose representatives $e_{o,i,o',1}, \dots, e_{o,i,o',j_{o'}}$ of the $G_{v_{o,i}}$ -orbits in $\delta^{-1}(o') \cap \text{lk}(v_{o,i})$. For each i and j choose

$$\Phi_{o,i,o',j} \in \text{QIsom}((X_{v_{o,i}}, \mathcal{P}_{v_{o,i}}, \delta, \xi \cdot \zeta), (X_{v_{o,1}}, \mathcal{P}_{v_{o,1}}, \delta, \xi \cdot \zeta))$$

such that $(\Phi_{o,i,o',j})_*(e_{o,i,o',j}) = e_{o,1,o',1}$. Such quasi-isometries exist by Proposition 6.9.

Similarly choose representatives $v'_{\beta(o),1}, \dots, v'_{\beta(o),i'_{\beta(o)}}$ of the G' -orbits contained in $(\delta')^{-1}(\beta(o))$ and representatives $e'_{\beta(o),i,\beta(o'),1}, \dots, e'_{\beta(o),i,\beta(o'),j'_{\beta(o')}} of the $G'_{v_{\beta(o),i}}$ -orbits in $(\delta')^{-1}(\beta(o')) \cap \text{lk}(v'_{\beta(o),i})$ and quasi-isometries $\Phi'_{\beta(o),i,\beta(o'),j}$.$

Choose a quasi-isometry $\Phi_{o,o'} \in \text{QIsom}((X_{v_{o,1}}, \mathcal{P}_{v_{o,1}}, \delta, \xi \cdot \zeta), (X'_{v'_{\beta(o),1}}, \mathcal{P}'_{v'_{\beta(o),1}}, \delta', \zeta'))$ that takes $e_{o,1,o',1}$ to $e'_{\beta(o),1,\beta(o'),1}$. Such a quasi-isometry exists by condition 3 and Proposition 6.9. If $e_{o,1,o',1}$ is ζ -unoriented then we also choose

$$\Phi_{o,o'}^- \in \text{QIsom}((X_{v_{o,1}}, \mathcal{P}_{v_{o,1}}, \delta, \xi \cdot \zeta), (X'_{v'_{\beta(o),1}}, \mathcal{P}'_{v'_{\beta(o),1}}, \delta', \zeta'))$$

that takes $e_{o,1,o',1}$ to $e'_{\beta(o),1,\beta(o'),1}$ such that $\Phi_{o,o'}^- \circ (\Phi_{o,o'})^{-1}$ orientation reversing on $X'_{e'_{\beta(o),1,\beta(o'),1}}$. Such a quasi-isometry exists by Corollary 6.10.

We have chosen finitely many quasi-isometries Φ , so they have uniformly bounded quasi-isometry constants.

Induction base case Begin the induction by choosing a cylindrical vertex $c \in T$ and a cylindrical vertex $c' \in (\delta')^{-1}(\beta(\delta(c)))$. Define $\chi(c) := c'$. By construction X_c and $X'_{c'}$ are copies of \mathbb{R} . Define $\phi_c: X_c \rightarrow X'_{c'}$ to be an isometry. If c is ζ -oriented we choose ϕ_c so that $\phi_c(\xi \cdot \zeta(X_c)) = \zeta'(X'_{c'})$. Extend χ to $\text{lk}(c)$ as in Theorem 6.13.

Inductive steps for non-elementary vertices Suppose $v = \tau(e)$ is a non-elementary vertex such that for $c = \iota(e)$ we have already defined a coarse isometry $\phi_c: X_c \rightarrow X'_{\chi(c)}$ and $\chi|_{\text{lk}(c)}$. Suppose further that if c is ζ -oriented then $\phi_c(\xi \cdot \zeta(X_c)) = \zeta'(X'_{\chi(c)})$. Let $c' := \chi(c)$ and $e' := \chi(e)$.

Suppose v is rigid. Now $gv = v_{\delta(v),i}$ for some $g \in G$ and some i , and $hge = e_{\delta(v),i,\delta(e),j}$ for some j and some $h \in G_{v_{\delta(v),i}}$. Similarly, there are $g' \in G'$ and i' such that $g'v' = v'_{\beta(\delta(v)),i'}$, and j' and $h' \in G'_{v'_{\beta(\delta(v)),i'}}$ such that $h'g'e' = e'_{\beta(\delta(v)),i',\beta(\delta(e)),j'}$.

The map

$$(h'g')^{-1} \circ (\Phi'_{\beta(\delta(v)),i',\beta(\delta(e)),j'})^{-1} \circ \Phi_{\delta(v),\delta(e)} \circ \Phi_{\delta(v),i,\delta(e),j} \circ hg$$

is an element of $\text{QIsom}(X_{v_0}, \mathcal{P}_{v_0}, \beta \circ \delta, \xi \cdot \zeta), (X'_{v'_0}, \mathcal{P}'_{v'_0}, \delta', \zeta')$ taking e to e' .

If e is ζ -unoriented then we also have that the map

$$(h'g')^{-1} \circ (\Phi'_{\beta(\delta(v_0)), i', \beta(\delta(e)), j'})^{-1} \circ \Phi_{\delta(v_0), \delta(e)}^{-1} \circ \Phi_{\delta(v_0), i, \delta(e), j} \circ hg$$

is an element of $\text{QIsom}(X_{v_0}, \mathcal{P}_{v_0}, \beta \circ \delta, \xi \cdot \zeta), (X'_{v'_0}, \mathcal{P}'_{v'_0}, \delta', \zeta')$ taking e to e' . For one of these two, the composition with $\alpha_e \circ (\phi_c)^{-1} \circ (\alpha'_{e'})^{-1}$ is orientation preserving on $X_{\bar{e}}$. Choose this one as ϕ'_v .

Finally, choose a point $x \in X_{\bar{e}}$. Since edge stabilizers act uniformly coboundedly on their corresponding peripheral sets, we can choose an element $k \in G'_{\bar{e}'}$ that is orientation preserving on $X'_{\bar{e}'}$ and such that $k\phi'_v(x)$ is boundedly close to $\alpha'_{e'} \circ \phi_c \circ (\alpha_e)^{-1}(x)$. Define $\phi_v := k\phi'_v$.

We have:

- $\phi_v \in \text{QIsom}((X_{v_0}, \mathcal{P}_{v_0}, \beta \circ \delta, \xi \cdot \zeta), (X'_{v'_0}, \mathcal{P}'_{v'_0}, \delta', \zeta'))$
- $(\phi_v)_*(e) = e'$
- $\phi_v(x)$ is boundedly close to $\alpha'_{e'} \circ \phi_c \circ (\alpha_e)^{-1}(x)$.
- $\phi_v \circ (\alpha'_{e'} \circ \phi_c \circ (\alpha_e)^{-1})^{-1}$ is orientation preserving on $X'_{\bar{e}'}$.
- ϕ_v is a composition of three of the pre-chosen Φ with multiplication by five group elements, so the quasi-isometry constants of ϕ_v are bounded in terms of those of the Φ and the constants for the group action.

By relative quasi-isometric rigidity, ϕ_v is a coarse isometry.

We also claim that $\alpha'_{e'} \circ \phi_c \circ (\alpha_e)^{-1}$ is a coarse isometry. This is because ϕ_c is a coarse isometry, by the induction hypothesis, and α_e and $\alpha'_{e'}$ are, by construction (recall Remark 8.2), coarse similitudes with multiplicative constants $\text{relStr}(e)$ and $\text{relStr}(e')$, which are equal, since:

$$\text{relStr}(e) = \delta_0 \circ \delta^{-1}(\delta(e)) = \delta'_0 \circ (\delta')^{-1} \circ \beta(\delta(e)) = \delta'_0 \circ (\delta')^{-1}(\delta'(e')) = \delta'_0(e') = \text{relStr}(e')$$

Thus, $\phi_v \circ (\alpha'_{e'} \circ \phi_c \circ (\alpha_e)^{-1})^{-1}$ is orientation preserving coarse isometry on $X'_{\bar{e}'}$ that coarsely fixes a point. It follows that $\phi_v|_{X_{\bar{e}}}$ and $\alpha'_{e'} \circ \phi_c \circ (\alpha_e)^{-1}$ are coarsely equivalent.

Define $\chi|_{\text{lk}(v)} := (\phi_v)_*$.

For each edge $e'' \in \text{lk}(v) \setminus \{\bar{e}\}$ define $\phi_{\tau(e'')} := \alpha'_{\chi(e'')} \circ \phi_v \circ (\alpha_{e''})^{-1}$. Since ϕ_v is a coarse isometry and $\alpha_{e''}$ and $\alpha'_{\chi(e'')}$ are coarse similitudes with the same multiplicative constant, as above, we have that $\phi_{\tau(e'')}$ is a coarse isometry.

Suppose v is hanging. The map $\alpha'_{e'} \circ \phi_c \circ (\alpha_e)^{-1}: X_{\bar{e}} \rightarrow X'_{\bar{e}'}$ is a coarse isometry, since attaching maps to hanging vertex spaces are coarse isometries by Lemma 8.3 and ϕ_c is a coarse isometry by the induction hypothesis.

Use condition 3 and Proposition 8.1 to produce a quasi-isometry

$$\phi_v \in \text{QIsom}((X_v, \mathcal{P}_v, \beta \circ \delta, \xi \cdot \zeta), (X'_{v'}, \mathcal{P}'_{v'}, \delta', \zeta'))$$

that is a coarse isometry along each peripheral subset and that coarsely agrees with $\alpha'_{e'} \circ \phi_c \circ (\alpha_e)^{-1}$ on $X_{\bar{e}}$.

Define $\chi|_{\text{lk}(v)} := (\phi_v)_*$.

For each $e'' \in \text{lk}(v) \setminus \{\bar{e}\}$ the map $\phi_{\tau(e'')} := \alpha'_{\chi(e'')} \circ \phi_v \circ (\alpha_{e''})^{-1}$ is a coarse isometry, since attaching maps to hanging vertex spaces are coarse isometries by Lemma 8.3, and ϕ_v is a coarse isometry along peripheral sets by construction.

Inductive step for cylindrical vertices Suppose $c = \iota(e)$ is cylindrical, χ is defined on e , $\phi_{\tau(e)}$ is defined, and ϕ_c is a coarse isometry such that ϕ_c is coarsely equivalent to $(\alpha'_{\chi(e)})^{-1} \circ \phi_{\tau(e)} \circ \alpha_e$. Extend χ to $\text{lk}(c) \setminus \{e\}$ as in Theorem 6.13.

This completes the induction. The result is $\chi \in \text{Isom}((T, \delta), (T', \delta'))$ and uniform quasi-isometries (ϕ_v) satisfying the conditions of Corollary 2.18, so (ϕ_v) is a tree of quasi-isometries over χ compatible with X and X' , as desired. \square

References

- [1] Jason A. Behrstock, Bruce Kleiner, Yair N. Minsky, and Lee Mosher, *Geometry and rigidity of mapping class groups*, Geom. Topol. **16** (2012), no. 2, 781–888.
- [2] Jason A. Behrstock and Walter D. Neumann, *Quasi-isometric classification of graph manifold groups*, Duke Math. J. **141** (2008), no. 2, 217–240.
- [3] Jason A. Behrstock and Walter D. Neumann, *Quasi-isometric classification of non-geometric 3-manifold groups*, J. Reine Angew. Math. **669** (2012), 101–120.
- [4] M. Bestvina and M. Feighn, *A combination theorem for negatively curved groups*, J. Differential Geom. **35** (1992), no. 1, 85–101.
- [5] Kingshook Biswas and Mahan Mj, *Pattern rigidity in hyperbolic spaces: duality and PD subgroups*, Groups Geom. Dyn. **6** (2012), no. 1, 97–123.
- [6] M. Bonk and O. Schramm, *Embeddings of Gromov hyperbolic spaces*, Geom. Funct. Anal. **10** (2000), no. 2, 266–306.
- [7] Marc Bourdon and Hervé Pajot, *Rigidity of quasi-isometries for some hyperbolic buildings*, Comment. Math. Helv. **75** (2000), no. 4, 701–736.
- [8] Marc Bourdon and Hervé Pajot, *Cohomologie l_p et espaces de Besov*, J. Reine Angew. Math. **558** (2003), 85–108.
- [9] Brian H. Bowditch, *Cut points and canonical splittings of hyperbolic groups*, Acta Math. **180** (1998), no. 2, 145–186.
- [10] Martin R. Bridson and André Haefliger, *Metric spaces of non-positive curvature*, Grundlehren der mathematischen Wissenschaften, vol. 319, Springer, Berlin, 1999.

- [11] Sergei Buyalo and Viktor Schroeder, *Elements of asymptotic geometry*, EMS Monographs in Mathematics, European Mathematical Society (EMS), Zürich, 2007.
- [12] Christopher H. Cashen, *Splitting line patterns in free groups*, Algebr. Geom. Topol., in press, [arXiv:1009.2492](#).
- [13] Christopher H. Cashen, *Quasi-isometries between tubular groups*, Groups Geom. Dyn. **4** (2010), no. 3, 473–516.
- [14] Christopher H. Cashen and Nataša Macura, *Line patterns in free groups*, Geom. Topol. **15** (2011), no. 3, 1419–1475.
- [15] Pallavi Dani and Anne Thomas, *Quasi-isometry classification of certain right-angled Coxeter groups*, preprint (2014), [arXiv:1402.6224v4](#).
- [16] Martin J. Dunwoody, *The accessibility of finitely presented groups*, Invent. Math. **81** (1985), no. 3, 449–457.
- [17] Martin J. Dunwoody and Michah E. Sageev, *JSJ-splittings for finitely presented groups over slender groups*, Invent. Math. **135** (1999), no. 1, 25–44.
- [18] Max Forester, *On uniqueness of JSJ decompositions of finitely generated groups*, Comment. Math. Helv. **78** (2003), no. 4, 740–751.
- [19] Koji Fujiwara and Panos Papasoglu, *JSJ-decompositions of finitely presented groups and complexes of groups*, Geom. Funct. Anal. **16** (2006), 70–125.
- [20] Mikhael Gromov, *Infinite groups as geometric objects*, Proceedings of the International Congress of Mathematicians, Vol. 1, 2 (Warsaw, 1983) (Warsaw), PWN, 1984, pp. 385–392.
- [21] Mikhael Gromov, *Hyperbolic groups*, Essays in group theory, Math. Sci. Res. Inst. Publ., vol. 8, Springer, New York, 1987, pp. 75–263.
- [22] Vincent Guirardel and Gilbert Levitt, *JSJ decompositions: definitions, existence, uniqueness. I: The JSJ deformation space*, preprint (2009), [arXiv:0911.3173](#).
- [23] Vincent Guirardel and Gilbert Levitt, *Trees of cylinders and canonical splittings*, Geom. Topol. **15** (2011), no. 2, 977–1012.
- [24] Ilya Kapovich, *The combination theorem and quasiconvexity*, Internat. J. Algebra Comput. **11** (2001), no. 2, 185–216.
- [25] Michael Kapovich and Bruce Kleiner, *Hyperbolic groups with low-dimensional boundary*, Ann. Sci. École Norm. Sup. (4) **33** (2000), no. 5, 647–669.
- [26] Bruce Kleiner and Bernhard Leeb, *Rigidity of quasi-isometries for symmetric spaces and Euclidean buildings*, Inst. Hautes Études Sci. Publ. Math. (1997), no. 86, 115–197 (1998).

- [27] Frank Thomson Leighton, *Finite common coverings of graphs*, J. Combin. Theory Ser. B **33** (1982), no. 3, 231–238.
- [28] William Malone, *Topics in geometric group theory*, Ph.D. thesis, University of Utah, Salt Lake City, UT, 2010.
- [29] Vladimir Markovic, *Quasisymmetric groups*, J. Amer. Math. Soc. **19** (2006), no. 3, 673–715.
- [30] Alexandre Martin, *Non-positively curved complexes of groups and boundaries*, Geom. Topol. **18** (2014), no. 1, 31–102.
- [31] Alexandre Martin and Jacek Świątkowski, *Infinitely-ended hyperbolic groups with homeomorphic Gromov boundaries*, J. Group Theory **18** (2015), no. 2, 273–289.
- [32] Mahan Mj, *Pattern rigidity and the Hilbert-Smith conjecture*, Geom. Topol. **16** (2012), no. 2, 1205–1246.
- [33] Lee Mosher, Michah Sageev, and Kevin Whyte, *Quasi-actions on trees II: Finite depth Bass-Serre trees*, Mem. Amer. Math. Soc. **214** (2011), no. 1008, vi+105.
- [34] Walter D. Neumann, *On Leighton’s graph covering theorem*, Groups Geom. Dyn. **4** (2010), no. 4, 863–872.
- [35] Pierre Pansu, *Métriques de Carnot-Carathéodory et quasiisométries des espaces symétriques de rang un*, Ann. of Math. (2) **129** (1989), no. 1, 1–60.
- [36] Panos Papasoglu, *Quasi-isometry invariance of group splittings*, Ann. of Math. (2) **161** (2005), no. 2, 759–830.
- [37] Panos Papasoglu and Kevin Whyte, *Quasi-isometries between groups with infinitely many ends*, Comment. Math. Helv. **77** (2002), no. 1, 133–144.
- [38] E. Rips and Z. Sela, *Structure and rigidity in hyperbolic groups. I*, Geom. Funct. Anal. **4** (1994), no. 3, 337–371.
- [39] E. Rips and Z. Sela, *Cyclic splittings of finitely presented groups and the canonical JSJ decomposition*, Ann. of Math. (2) **146** (1997), no. 1, 53–109.
- [40] Richard Evan Schwartz, *Symmetric patterns of geodesics and automorphisms of surface groups*, Invent. Math. **128** (1997), no. 1, 177–199.
- [41] Jean-Pierre Serre, *Trees*, Springer Monographs in Mathematics, Springer-Verlag, Berlin, 2003, Translated from the French original by John Stillwell, Corrected 2nd printing of the 1980 English translation.
- [42] John R. Stallings, *On torsion-free groups with infinitely many ends*, Ann. of Math. (2) **88** (1968), 312–334.

- [43] John R. Stallings, *Group theory and three-dimensional manifolds*, Yale University Press, New Haven, Conn., 1971, A James K. Whittemore Lecture in Mathematics given at Yale University, 1969, Yale Mathematical Monographs, 4.
- [44] Pekka Tukia, *Quasiconformal extension of quasisymmetric mappings compatible with a Möbius group*, Acta Math. **154** (1985), no. 3-4, 153–193.
- [45] Diane M. Vavrichek, *The quasi-isometry invariance of commensurizer subgroups*, Groups Geom. Dyn. **7** (2013), no. 1, 205–261.
- [46] Xiangdong Xie, *Quasi-isometric rigidity of Fuchsian buildings*, Topology **45** (2006), no. 1, 101–169.
- [47] Xiangdong Xie, *Quasisymmetric maps on the boundary of a negatively curved solvable Lie group*, Math. Ann. **353** (2012), no. 3, 727–746.

FAKULTÄT FÜR MATHEMATIK, UNIVERSITÄT WIEN, OSKAR-MORGENSTERN-PLATZ 1, 1090 WIEN,
 ÖSTERREICH
 christopher.cashen@univie.ac.at
 alexandre.martin@univie.ac.at